

Over time RF fitting for Jitter Free 3D Vertebra Reconstruction from Video Fluoroscopy

Ioannis Ioannidis and Hammadi Nait-Charif

Bournemouth University, Poole BH12 5BB, UK
yioannidis@bournemouth.ac.uk

Abstract. Over the past decades, there has been an increasing interest in spine kinematics and various approaches have been proposed on how to analyse spine kinematics. Amongst all, emphasis has been given to both the shape of the individual vertebrae as well as the overall spine curvature as a mean of providing accurate and valid spinal condition diagnosis. Traditional invasive methods cannot accurately delineate the intersegmental motion of the spine vertebrae. On the contrary, capturing and measuring spinal motion via the non-invasive fluoroscopy has been a popular technique choice because of its low incurred patient radiation exposure nature. In general, image-based 3D reconstruction methods focus on static spine instances. However, even the ones analysing sequences yield in unstable and jittery animations of the reconstructed spine. In this paper, we address this issue using a novel approach to robustly reconstruct and rigidly derive a shape with no inter-frame variations. This is to produce animations that are jitter free across our sequence based on fluoroscopy video. Our main contributions are 1) retaining the shape of the solid vertebrae across the frame range, 2) helping towards a more accurate image segmentation even when there's a limited training set. We show our pipeline's success by reconstructing and comparing 3D animations of the lumbar spine from a corresponding fluoroscopic video.

1 Introduction

In orthopaedics, identifying and accurately reconstructing 3D spine has always been of a great interest. Accurate shape and motion of the spine is essential to both physicians, medical students and patients. It's important because it's the means of understanding, assessing and providing a spinal condition diagnosis. Spine shape and kinematics reconstruction has been extensively studied [27, 28]

Many studies focus on individual vertebrae surface reconstruction or use pre-defined rigid models (usually from CT scans) for the quantification of in vivo inter-vertebral motion. They are fitted to sequences of images and roto-scoped and registered on the target frames or on each subsequent frame in video sequences [27, 28]. From an animation point of view, these approaches are limited in a sense that potential video sequences would be reconstructed statically per frame. Hence the per frame errors would result in jittery and unstable reconstructed animations. These methods lack shape representation as they assume

the vertebrae shapes stay solid throughout the captured image sequence, not accounting for projection noise across a sequence. This is not always the case, even though spine vertebrae are theoretically solid, and it's mostly because of the presence of noise in the target images. To address issues such as incorrect model initialization, segmentation and tracking error accumulation, a number of successful techniques have been developed to minimize the resulting errors for both surface reconstruction and kinematic analysis [19, 3]. The majority of them are solving temporal coherency based on some sort of interpolation method as in [27]. Many, such as Prabhu et al. [18], have used Kalman Filtering (KF) to smooth the fitted landmarks after fitting. Others used either the current shape parameters and (KF) to update the view space per frame or even used convolutional neural networks (CNN) to accurately track and reconstruct [10, 17].

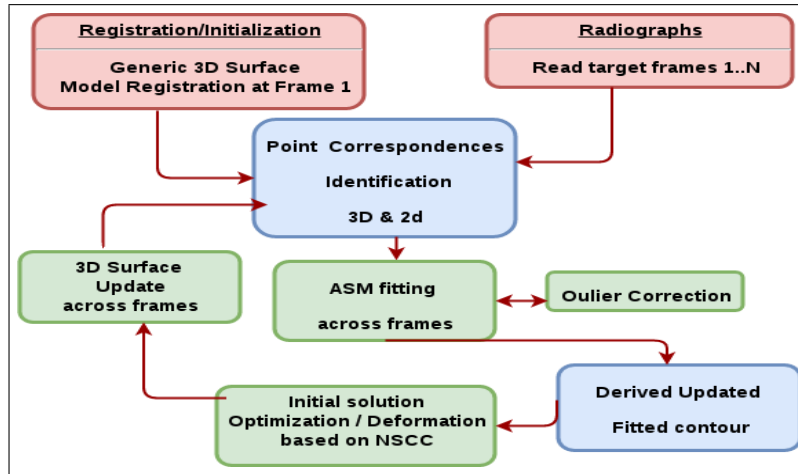


Fig. 1. Animation Reconstruction Pipeline

Crivellaro et al. [8] estimate the pose in a given frame, and then add it as a component of the pose prior for the next frame using KF for smoothing. Liao et al. [14] apply 3D score maps on a per CT slice basis (for the same frame). Hence leaving out the over time segmentation information whereas [2, 1] work on neighbour slices only and consecutive frames. Wand et al. [24] achieve the global smoothness over time by iteratively merging geometry from adjacent frames. The work in [12] is based on as rigid as possible (ARAP) principle presented in [21]. It assumes there's a slight deformation model that is guided from frame to frame. In [26] CNN are employed however as it's obvious from the results, the final tracking is jittery. Salzmann et al. [20] smooth out predictions over a predefined set of frames to alleviate jitter.

Despite the above mentioned methods applied time varying analysis filters to smooth out and improve tracking accuracy, they usually suffer from 2 major problems. At first, they are sensitive to outliers. Secondly, inevitably the smoothing applied amongst the shapes between frames derives slightly different per frame morphologies for the final tracked objects. This is unwanted, as

the aim is to rigidly segment from pose to pose and across all frames based on the underlying image information so as to produce a realistic animation with minimal jitter. Our method addresses the above by relying on random forests (RF) for the selection of displacements for all shapes' correction across the frame range simultaneously. More analytically, we used the whole target frame range, such that each next iteration prediction, for each shape of the sequence, is based on a collection of RF predictions across the whole frame range from the previous iteration. This way we can improve segmentation, as wrongly chosen individual landmark displacements obey a global per iteration displacement factor. Thus they can't severely affect the next frame segmentation in case of accumulated errors. Possible shape deviations are corrected as part of a more general drag towards the most significant displacements imposed to the rest of the shapes across the frame range. In our method we applied RF over time across the test frame range to fit the reconstructed 3D surface model to our video sequence. This is a different approach compared to [15] where RF have been used to improve individual image segmentation.

In our method, we used as a basis and extended the traditional active shape model (ASM) [6]. ASM has been extensively used mostly because of it's proven robustness especially when dealing with large image inconsistencies and noise variations which is what fluoroscopy produces. The main contributions of this paper are:

- A novel approach which uses ASM allowing it to retain the shape of the supposedly solid vertebrae across the frame range
- A more accurate image segmentation, as a consequence of the first contribution, even when there's a limited training set because of global drag towards the most significant displacements

2 Method Overview

In Figure 1 we present the pipeline of our method. We start with the input of both target images and a generic 3D surface model for a particular vertebra. These are then fed to the point correspondence step where we identify the 3D model points that correspond to the 2D radiographic landmarks. Then, ASM fitting is performed across the target frame range simultaneously for all frames using RF for the choice displacement predictions and the derived new contours are computed for each frame at this iteration. After having acquired the updated contours, we update our initial surface model solution with the derived contours and deform across all target frames for this iteration based on an adaptation of the non stereo corresponding contour (NSCC) algorithm [13] and the cycle continues.

3 Methods

In this section, we analytically present our method of extending ASM by iteratively fitting over time using RF voting (OTRF).

3.1 Traditional ASM method

In ASM [6], a statistical shape model (SSM) framework, also named point distribution model (PDM) is used to fit a shape contour to a 2D image by iteratively deforming the initial shape instance according to the PDM. Each shape is a vector of 2D landmark points $v_i = \{l_{i_1}..l_{i_n}\}$, where n is the total number of points of the shape. For each shape in the training set a v vector is created and all vectors are aligned via principle component analysis (PCA) [6]. Shape alignment results in a mean shape $\bar{x} = \frac{1}{n} \sum_{i=1}^n v_i$, where n is the total number of training shapes. Now, after PCA we can describe our model with the following:

$$x \approx \bar{x} + Pb \quad (1)$$

,where b is a vector of the most significant shape coordinate parameters of the deformable model [6] and P is the covariance matrix of first t eigenvectors calculated during PCA. This matrix contains eigenvectors and their corresponding eigenvalues λ_i . The largest eigenvalue describes the most significant shape variation and only a few t eigenvalues are needed to describe most of the shape's variation [4]. Equation 1 is then re-written to the following:

$$b = P^{-1}(x_i - \bar{x}) \quad (2)$$

By varying the parameters of b within $-3\lambda_i < b_i < 3\lambda_i$, we can vary the mean shape towards the most significant eigenvectors of P matrix when fitting to a target image. So, at first during training, for each of the landmarks of the shape, we create SSM derivative profiles (see Figure 2) sampled along the normal of each landmark across the frame range of images, $dg_{i_j} = [g_{i_{j_1}} - g_{i_{j_0}}, \dots, g_{i_{j_{2k+1}}} - g_{i_{j_{2k}}}]$, where dg_{i_j} is each landmark's derivative profile, k is the number of sampled profile points on either side of the landmark along the normal, i refers to every image of the training set and j to each landmark of the shape across images [4]. Their corresponding covariance matrices C are also computed

$$C_i = \frac{1}{n} \sum_{i=1}^n (g_{i_j} - \bar{g}_j)(g_{i_j} - \bar{g}_j)^T \quad (3)$$

and we finally derive the mean normalized derivative profile for each landmark across the frame range.

$$\bar{g}_j = \frac{1}{n} \sum_{i=1}^n g_{i_j} \quad (4)$$

, where n is the total number of images in the training set. Then during fitting, for each of the landmarks, we create a derivative search profile by sampling small segments of m pixels across n pixels on either side of our landmark along the normal ($m < n$) and compare to the mean normalized derivative profile using Mahalanobis distance,

$$D = (g_j(d) - \bar{g}_j)S^{-1}(g_j(d) - \bar{g}_j) \quad (5)$$

, where g_j is current search profile, \bar{g} is the mean model profile, d is the center pixel of the current profile sample across the normal and S is the mean covariance matrix for this specific landmark. Minimizing D is equal to maximizing the probability of $g(d)$, hence the d that achieves the minimum D is the one towards which this landmark should be moved [4]. We do this for all landmarks to derive their new most likely locations on the image. Finally, at every fitting iteration we update the b model parameters with $g_j(d)$

$$b = P^T(g_j(d) - \bar{x}) \quad (6)$$




Fig. 2. Search profile along landmark normal

3.2 Proposed Method

In our approach, we extend ASM by iteratively fitting across the test frame range simultaneously using random forests (RF) for the choice displacement predictions. We refer to it as 'over time' prediction as well. The traditionally derived new contour predictions are computed for each frame at each iteration and are used to simultaneously feed each next iteration's shape fitting parameters, across the frame ranges tested. More analytically, at each iteration i we build displacement choice decision trees $T_{j_{i,r}}$ from each test frame range r for each of the landmarks of our shape.

$$T_{j_{i,r}} = \{T_{i_{j_1}} .. T_{i_{j_r}}\} \quad (7)$$

, where j is the current landmark, i is the current iteration and r total number of frame ranges to simultaneously segment. Consequently, we use Equation 7 across the shape landmarks

$$S = \{T_1 .. T_n\} \quad (8)$$

, where n is the total number of landmarks and i and j subscripts have been omitted in Equation 8 for notation simplicity. Once, our trees are built per iteration, we use them to build and train a random forest and then derive a prediction per landmark to represent the most likely best choice for displacement for this landmark across the test frame ranges. From S we derive a vector of our shape predictions, for each landmark separately

$$fp_i = \{fp_{i_1} .. fp_{i_j}\} \quad (9)$$

, where j represents a single landmark of our shape and i is the current iteration number.

In Figure 3 we show how we used the different searched frame ranges across the target frames to produce trees, as well as how each of them is fed to the RF so as to finally produce a displacement vector per iteration out of the predictions per landmark across these ranges. The idea is that by searching over time, and not on a per frame basis, each next frame’s segmentation doesn’t rely solely on static image information. Similarly, the pose across the target frame ranges doesn’t depend exclusively on the each previous frame’s pose. In the traditional ASM, during search phase we scan along the normal to the point in question or it’s surrounding area and use a metric such as Mahalanobis distance to express the likelihood of the current searched profile compared to the training set’s equivalent profile as shown in Figure 2, per each individual frame.

In our method, we harvest the underlying image information for both shape and pose over time and across all target frames and compute a refined shape that is insensitive to local noise and takes into account both local and global shape variations. Because of this, the shape keeps it’s rigidity across the sequence and overcomes the possibly inaccurate smoothing caused by interpolating 2 frames (as KF based approaches [7]). For example, in this paper Cordea et al. [7] used an extended Kalman filtering approach and applied it onto Active appearance models (AAM) [5]. They updated the shape and pose parameters per iteration during fitting. Huang et al. here [10] also used an this approach to update shape and pose parameters and refine per frame, along with a temporal matching filter to smooth inter-frame shape differences.

These methods inevitably result in some sort of smoothing because of the nature of Kalman filtering itself. In medical imaging and specifically fluoroscopy, a low-radiation incurring technique, there is a high possibility of increased amounts of variation even between successive frames. Hence, one frame might have been precisely segmented and another not. This would result in an in between displacement for this particular landmark which would deviate from the image feature. The smoothed interpolated position that KF approaches are based on, would then result in successive frames to be clearly different from each other.

In our method, what we are suggesting is that the final shape is as solid as possible (subject to sub pixel inaccuracy), driven by the information of the whole range of target images and any inter-frame shape inconsistencies are only due to sub pixel inaccuracies caused during segmentation. Thus, producing a smooth and overall accurate animation result with no visually distracting jitter. We manage to achieve this using over time RF (OTRF) displacement choices across our frame ranges as described previously.

Most KF based and extended KF (EKF) approaches that update the shape are either updated within each frame [10] using information from previous iterations, between successive frames [17] or across a sequence to reinforce shape parameters [7]. In [25] KF is used to correct the global rigid motion after fitting has converged on each frame. Consequently, the use of KF which is mostly used to derive a better initial shape for each of the next frames, is not involved in

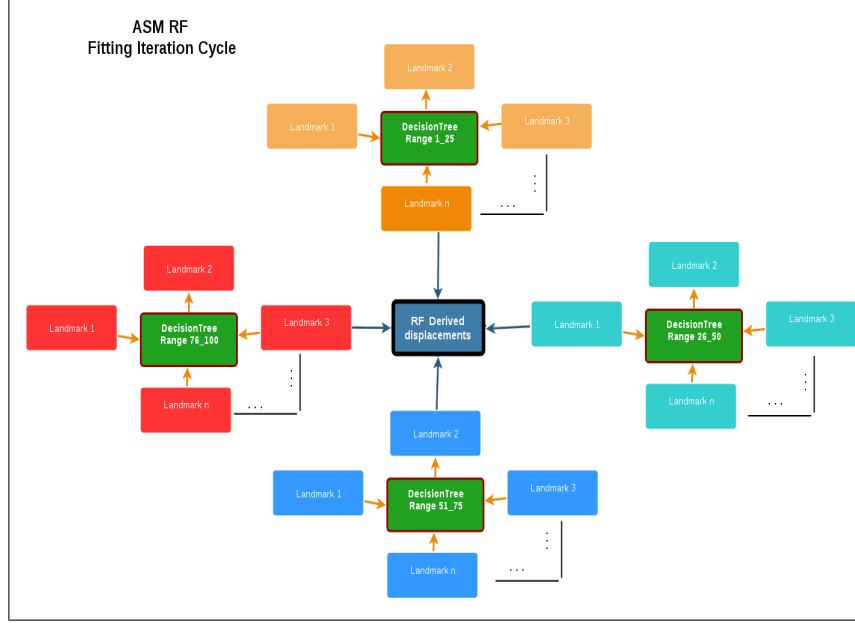


Fig. 3. RF fitting across frame ranges

iterative shape fitting. Additionally, KF is usually accompanied with a smoother or similar method that smooths the in between observations. This is because KF methods are generally sensitive to outliers and there is a need for a more robust outlier detection method to handle tracking inaccuracies. Wadehn et al. show in [23], a smoother that was put in place to address possible outliers. Others, such as Ting et al. in [22] followed a weighted least squares approach by applying weights to data observations to tackle outliers and hence improve tracking robustness.

However, all the above mentioned methods could result in some sort of inter-frame inaccuracies owed to the nature of KF and consequently shape inconsistencies in an animation reconstruction pipeline. The shape will change and adapt to inter-frame smoothing and in the presence of outliers successive frame shapes will vary even more significantly. This is undesirable when our aim is to robustly reconstruct and rigidly derive a shape with no inter-frame variations and produce animations that are jitter free across our sequence.

3.3 Reconstruction

The rest of the reconstruction pipeline is mostly based on an adaptation of the NSCC algorithm [13], which practically updates the surface of a 3D model from a contour based on correspondences between associated 3D surface and 2D contour points. In this paper we use a single plane fluoroscopy video to segment lumbar spine contours, assuming that there's no in-plane rotation. In the future




Fig. 4. *Maya reconstruction pipeline. From 2D segmented points imported (up left image), all the way down to final model deformation (bottom right image)*

we are planning to adapt our method using bi-planar sequences so as to achieve a more robust shape representation. As part of ASM we trained our model on a 2D fluoroscopy sequence with the help of a radiographer. For each instance of our shape we used 43 landmarks to represent it. We have only focused on the 3D reconstruction of the vertebral body. The complete Maya reconstruction pipeline described in the following paragraph is depicted in Figure 4. At first, we identify point correspondence between our 2D contours and our 3D generic surface model within Autodesk Maya 3D package. Then, we bring into Maya the 2D segmented contours C_{2D} and superimpose them on the corresponding vertebra at a particular frame f (Figure 4, image 1). Secondly, we import our generic model M_{3D} and identify in planar view the corresponding 2D points on its 3D surface (Figure 4, image 2). Then, we align the model on the 2D contour on the radiograph (Figure 4, image 3) and after this the 3D projected contours to the model (hence the 3D projected to the 2D contour) (Figure 4, image 4). Finally, we iteratively deform the 3D projected contour C_{3D} to C_{2D} on this particular frame (Figure 4, image 5) and then we deform the rest of the M_{3D} to reflect the displaced C_{3D} (Figure 4, image 6). This way, the 3D model adapts to each new frame segmented with ASM and when this process is over we derive an animation based on the whole sequence segmentation. It is important to highlight that as our method is based on a 'model to images' correspondence to derive the 3D reconstruction, it's evident that the final reconstructed shape will be in immediate relation with the 2D segmentation. This is why it's crucial to derive a "solid" 2D shape in the first place.

When the reconstruction is completed across the target frame range we can monitor the inter-frame shape differences. By applying our method, we derive an almost rigid (subject to sub pixel inaccuracy) shape across the sequence whilst tracking the vertebra in question. This is owed to the fact that the 3D model shape is driven by and adapted exclusively to the 2D segmented contour. This segmented contour from ASM is iteratively updated based on the most RF voted

displacement and it remains solid. In contrast, the commonly used KF with a fixed lag smoother is causing a smoothing effect between observations.

Figure 5 shows the process of RF voting and choosing the most voted displacement. We are doing the same for all shape’s landmarks and derive a solid contour across frames. In Figure 6, the 2D contour drives the 3D shape at each frame of our video sequence and in Figure 7 we present the final reconstructed frames.




Fig. 5. *RF range voting across frames' at every iteration*

3.4 Experiments

Figure 7 shows 2 final reconstructed frames f_1, f_{75} . As seen in both Figure 6 (bottom right image 8) and Figure 7, there is no visible inter-frame surface variation amongst the compared frames whilst tracking is correctly achieved.

We implemented the traditional ASM method and a commonly used fixed-lag KF alongside with our over time RF fitting (OTRF) extension. Table 1 shows the quantified results of their mean inter-frame rmse (root mean squared error). We derive our final rmse metric f_c .

$$f_c = \frac{1}{n-1} \sum_{i=1}^{n-1} err_{i_{i+1}} \quad (10)$$

,where f_c is the inter-frame error metric across the sequence and i is every frame in our sequence. Finally, we can derive the average rmse for each method by averaging f_c across the frame range, as shown in Equation 11.

$$\bar{f}_c = \frac{1}{n} \sum_{i=1}^n f_c \quad (11)$$

As shown in Table 1, our proposed OTRF produces less inter-frame mean error than the traditional ASM and the KF based ASM and thus leading to

smoother and more accurate animation reconstructions which becomes more obvious in the presence of outliers or a limited training set.

Method	Inter-frame surface RMSE Sum - \bar{f}_c
ASM	1.2330
KF ASM	0.6891
OTRF ASM	0.6617

Table 1. *Inter-frame RMSE results across 75 frames.*

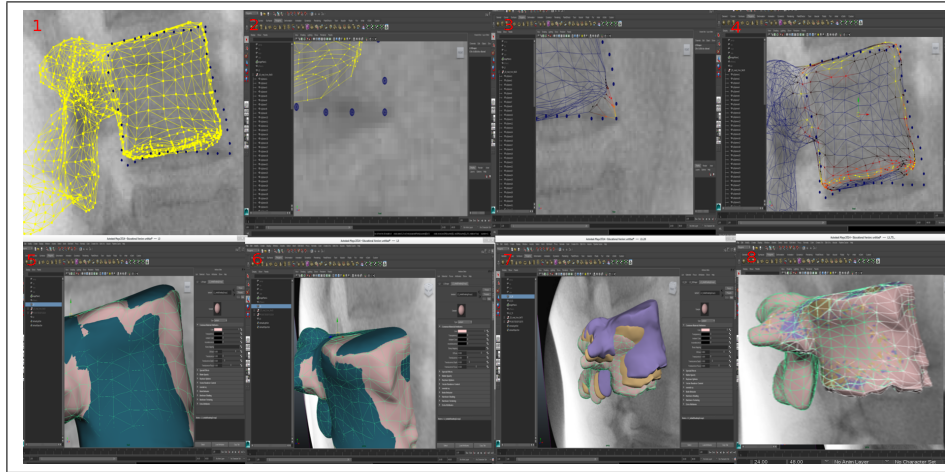


Fig. 6. *Reconstruction pipeline and final coloured difference comparison. Images 1 to 4: Individual 3D projected contour fitted to 2D segmented contour. Images 5 and 6: Frame f_{26}, f_{75} after fitting and reconstructed with KF and Traditional ASM. Images 7 and 8: f_{26}, f_{51}, f_{75} reconstruction with over time RF voting (OTRF) displacement ASM fitting. Image 8 superimposes OTRF reconstructed f_{26}, f_{51}, f_{75} to highlight the minimal colour differences across the range.*

3.5 Conclusion

In this paper, we introduced a novel extension method applied on ASM which addresses the issue of jittery 3D reconstruction animations when reconstructing from 2D images. We trained and segmented using a test fluoroscopy video of a lumbar spine. The reconstructed 3D shape retains its rigidity across the sequence and it doesn't suffer from inaccurate smoothing caused by interpolating adjacent frames. We used random forests as a voting method to select the next iteration's most likely displacements during ASM fitting, simultaneously across all frames. The algorithm outperformed the traditional ASM and KF implementations on the same dataset in terms of rigidity, whilst correctly segmenting the 2D and thus tracking the 3D vertebra in question. Our future plan is to improve the developed pipeline and extend it to incorporate another view so as to reconstruct using a biplane rather than a single plane set up.




Fig. 7. *final reconstructed frames f_1, f_{75}*

References

1. M. Allan, S. Ourselin, D. J. Hawkes, J. D. Kelly, D. Stoyanov. 3-d pose estimation of articulated instruments in robotic minimally invasive surgery. *IEEE Trans on Medical Imaging*, 37(5):1204–1213, 2018.
2. S. Wang, ZP. Liang. Regularized shape deformation for image segmentation. In 2001 IEEE Int Conf on Acoustics, Speech, and Signal Processing. Proceed, volume 3, pages 1569–1572, 2001.
3. T. Bassani, C. Ottardi, F. Costa, M. Brayda-Bruno, H.-J. Wilke, F. Galbusera. Semiautomated 3d spine reconstruction from biplanar radiographic images: Prediction of intervertebral loading in scoliotic subjects. *Frontiers in Bioengineering and Biotechnology*, 2017.
4. Y.-W. Chen and L. C. Jain. *Subspace Methods for Pattern Recognition in Intelligent Environment*. Springer, 2016.
5. T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *Computer Vision — ECCV*, pages 484–498, 1998. Springer.
6. T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active shape models - their training and application. *Comput. Vis. Image Underst.*, 61(1):38–59, Jan. 1995.
7. M. D. Cordea, E. M. Petriu, and D. C. Petriu. Three-dimensional head tracking and facial expression recovery using an anthropometric muscle-based active appearance model. *IEEE Trans on Instrumentation and Measurement*, 57(8):1578–1588, 2008.
8. A. Crivellaro, M. Rad, Y. Verdie, K. M. Yi, P. Fua, and V. Lepetit. Robust 3d object tracking from monocular images using stable parts. *IEEE Trans on Pattern Analysis and Machine Intelligence*, 40(6):1465–1479, 2018.
9. H. Gray, L. Hinkler Books Pty, H. Carter, and G. Davidson. *Gray’s Anatomy: The Classic 1860 Edition*. Hinkler Books, 2008.
10. C. Huang, X. Ding, and C. Fang. Pose robust face tracking by combining view-based aams and temporal filters. *Computer Vision and Image Understanding*, 116(7):777 – 792, 2012.
11. K. Ito, K. Hosoda, et al. Direct assessment of 3d foot bone kinematics using biplanar x-ray fluoroscopy and an automatic model registration method. *Journal of Foot and Ankle Research*, 2015.

12. M. Jaimez, T. J. Cashman, A. Fitzgibbon, J. Gonzalez-Jimenez, and D. Cremers. An efficient background term for 3d reconstruction and tracking with smooth surface models. In 2017 IEEE (CVPR), pages 2575–2583, 2017.
13. S. Laporte, W. Skalli, J. D. Guise, F. Lavaste, and D. Mitton. A biplanar reconstruction method based on 2d and 3d contours: Application to the distal femur. *Computer Methods in Biomechanics and Biomedical Engineering*, 6(1):1–6, 2003.
14. H. Liao, A. Mesfin, and J. Luo. Joint vertebrae identification and localization in spinal ct images by combining short- and long-range contextual information. *IEEE Transon Medical Imaging*, 37(5):1266–1275, 2018.
15. C. Lindner, P. A. Bromiley, M. C. Ionita, and T. F. Cootes. Robust and accurate shape model matching using random forest regression-voting. *IEEE Trans on Pattern Analysis and Machine Intelligence*, 37(9):1862–1874, 2015.
16. H. Nait-Charif, A. Breen, and P. Thompson. Vertebrae tracking in lumbar spinal video-fluoroscopy using particle filters with semi-automatic initialisation. In *Advances in Visual Computing*, pages 61–69, 2012. Springer.
17. N. Peterfreund. Robust tracking of position and velocity with kalman snakes. *IEEE Trans on Pattern Analysis and Machine Intelligence*, 21(6):564–569, 1999.
18. U. Prabhu, K. Seshadri, and M. Savvides. Automatic facial landmark tracking in video sequences using kalman filter assisted active shape models. In K. N. Kutulakos, editor, *Trends and Topics in Computer Vision*, pages 86–99, 2012, Springer.
19. Z. Rui, W. Huai-yu, and W. Ruo-hong. Facial feature point tracking algorithm based on eyes tracking and active appearance model step fitting. In 2015 34th Chinese Control Conf. (CCC), pages 3695–3701, 2015.
20. M. Salzmann, J. Pilet, S. Ilic, and P. Fua. Surface deformation models for nonrigid 3d shape recovery. *IEEE Trans on Pattern Analysis and Machine Intelligence*, 29(8):1481–1487, 2007.
21. O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing*, pages 109–116, Switzerland, Eurographics.
22. J. A. Ting, E. Theodorou, and S. Schaal. Learning an outlier-robust kalman filter. In J. N. Kok, J. Koronacki, R. L. d. Mantaras, S. Matwin, D. Mladenić, and A. Skowron, editors, : *ECML 2007*, pages 748–756, Springer .
23. F. Wadehn, L. Bruderer, J. Dauwels, V. Sahdeva, H. Yu, and H. Loeliger. Outlier-insensitive kalman smoothing and marginal message passing. In 2016 24th (EU-SIPCO), pages 1242–1246, 2016.
24. M. Wand, P. Jenke, Q. Huang, M. Bokeloh, L. Guibas, and A. Schilling. Reconstruction of deforming geometry from time-varying point clouds. In *Proceedings of the Fifth Eurographics Symp on Geometry Processing*, pages 49–58, Switzerland, 2007. Eurographics.
25. C. Wang and X. Song. Tracking facial feature points with prediction-assisted view-based active shape model. In *Face and Gesture 2011*, pages 259–264, 2011.
26. P. Wang, V. M. Patel, and I. Hacihaliloglu. Simultaneous segmentation and classification of bone surfaces from ultrasound using a multi-feature guided cnn. In A. F. Frangi, J. A. Schnabel, C. Davatzikos, C. Alberola-López, and G. Fichtinger, editors, *Med. Image Computing and Computer Assisted Intervention – MICCAI 2018*, pages 134–142, Cham, 2018. Springer.
27. W. Yu and G. Zheng. Atlas-Based 3D Intensity Volume Reconstruction from 2D Long Leg Standing X-Rays: Application to Hard and Soft Tissues in Lower Extremity, pages 105–112. Springer, 2018.
28. G. Zheng. Statistical Shape Models and Atlases: Application to 2D-3D Reconstruction in THA, pages 183–190. Springer, 2018.