FACULTY OF ENGINEERING – UNIVERSITY OF PORTO

Three-Dimensional Biplanar Reconstruction of the Scoliotic Spine for Standard Clinical Setup

Daniel Cardoso de Moura

Thesis submitted to the Department of Informatics Engineering in fulfilment of the requirements for the degree of Doctor of Philosophy in Informatics Engineering, supervised by Professor Jorge Manuel Gomes Barbosa, Department of Informatics Engineering, and Professor João Manuel Ribeiro da Silva Tavares, Department of Mechanical Engineering.

Porto, September 2010

Members of the jury

Professor Eugénio da Costa Oliveira (Universidade do Porto, chairman)
Professor Jorge Gomes Barbosa (Universidade do Porto, supervisor)
Professor João Manuel Ribeiro da Silva Tavares (Universidade do Porto, supervisor)
Professor Farida Cheriet (École Polytechnique Montréal)
Professor Jorge dos Santos Salvador Marques (Universidade Técnica de Lisboa)
Professor Miguel Fernando Paiva Velhote Correia (Universidade do Porto)
Professor Maria Cristina de Carvalho Alves Ribeiro (Universidade do Porto)

Ao meu Pai, à minha Mãe, ao meu Irmão, e à minha Bandeirinha

Acknowledgements

I wish to express sincere gratitude to Professor Jorge Barbosa, my supervisor, for my scientific guidance and constant support. Thank you Jorge for being always there, for your sincerity, for your patience, for the several discussions we had, and for sharing your knowledge with me. If it was not for you, Professor Miguel Correia and Professor João Tavares I would never be working on this subject that I found so appealing and challenging. Thank you Professor João Tavares for your guidance, for being available every time that I came to you, for the promptest replies I ever known of, for your thorough reviews, and for sharing your experience with me.

I would like to acknowledge everyone in Polytechnique Montréal that was involved in this last year and a half. Thank you so much Professor Farida Cheriet for opening the doors of your lab to me, for creating the link with Jonathan, for visiting us this year, for your precious advices, for sharing data with us, for making this collaboration work and evolve, and for your insight of our work. I am really grateful that our paths have crossed. Thank you Jonathan Boisvert, you are a fantastic colleague. I cannot thank you enough for sharing the statistical data with me, for the several conversations that we had, for being always available to collaborate, and for the knowledge you shared with me. I hope that we continue to work together. Thank you Lama Seoul, my pivot at LIV4D, Montreal. Always helpful and available, and a great person. Thank you Hassina for being a volunteer on the validation study. In general, thanks to everyone at Montreal who made my short stay there so pleasant and made this collaboration work.

I would like to express my gratitude to the clinical team in Porto that made this work possible. Thank you Doctor Ana Mafalda Reis for being always available, for your interest in the project, and for making the collaboration with Serviço Médico de Imagem Computarizada (SMIC) possible. Thank you Doctor Eusébio Gomes and Doctor Manuel Laranjeira for your interest in the project, for your

will to innovate in this field, and for being always available. Thank you SMIC for making available the imaging equipment as well as your personnel. Finally, a big thank you to the radiologic technologists of SMIC, which were always interested, good-humoured, and ready to help, despite the constant bustle of the radiology service.

Thank you Professor Eugénio Oliveira and Professor Augusto Sousa for your constant support and encouragement and for being so understandable in that which was the worst period of my life during the first year of my PhD. I would also like to thank my colleagues at LIACC, INEB, LOME and DEI that always provided a good environment either to work or to chat. Finally, thank you so much Marisa Martins for translating the abstract to French.

I also acknowledge Fundação para a Ciência e a Tecnologia for granting this doctoral degree (grant reference SFRH/BD/31449/2006), and Fundação Calouste Gulbenkian for granting my visit to École Polytechnique Montréal.

Last, but by no means the least, I would like to express my gratitude to my family and friends, for supporting me, encouraging me, and for your love. Thank you father, mother and brother for your love, pride, support and for always believing in me. Thank you Lúcia's family, in particular Mrs. Soledade for you important support that provided me valuable time and for your love. Thank you Lúcia, you are my cheer, my joy, my muse. Thank you for all the patience you had all these years, for taking care of me, for always believing and being there for me, in the bad and good moments.

Abstract

Scoliosis is characterized by a three-dimensional (3D) deformation of the spine that requires a 3D evaluation. However, conventional 3D imaging techniques have shown to be inadequate and, therefore, 3D reconstructions are typically performed from planar radiographs, one on the frontal plane and another on the lateral. Yet, this approach presents several challenges. The first is concerned with extracting 3D data from 2D radiographs. Currently, this is solved by calibration methods that make use of radiopaque objects and rotatory platforms that minimise patient positioning errors. However, these methods require considerable changes in radiological setups and protocols, besides introducing artefacts on the images that, sometimes, overlap anatomical structures of interest. After calibrating the acquired radiographs another problem arises, recovering the shape of the spine. The gold-standard methods require an extensive set of landmarks that must be manually identified on each radiograph, making them resource-consuming and error-prone. For this and the aforementioned reasons, the methods that are currently used for 3D reconstructing the spine are not viable for routine clinical evaluations on radiological services with standard setups.

In this thesis it was hypothesised that it is possible to build personalised geometrical models of scoliotic spines from two radiographs acquired in standard clinical environments by means of computational methods that require limited user-interaction. For accomplishing this goal, two methods are proposed here: a calibration method that makes use of a distance, easily measured on site, that aims at minimising the need for calibration objects; and a 3D reconstruction method based on the deformation of an articulated model of the spine that aims at providing fast and accurate 3D reconstructions.

In this thesis, it is shown by the first time that it is possible to recover the scale of the spine without calibration objects, although introducing a small calibration object improves both accuracy and robustness of the 3D reconstruc-

tions. In addition, it is shown that the proposed model enables to improve results of calibrations of biplanar radiographs based on the minimisation of the retro-projection error. The calibration method proposed here only requires minimal changes to the radiographs acquisition protocols, resulting in minimal artefacts on the radiographs. Thus, we conclude that it satisfies the requirements for being used in standard clinical environment.

Concerning recovering the shape of the spine, the proposed method achieved the fastest reconstruction times ever, including user-interaction time and computation time. Additionally, it achieved, within semi-supervised methods, the highest accuracy determining the location of vertebrae and the centres of their endplates for both mild and severe scoliotic patients. Furthermore, it was shown that the method enables to reliably calculate clinical indices, even by non-expert users. Therefore, we conclude that the 3D reconstruction method can be used for routine clinical evaluations.

Resumo

A escoliose é caracterizada por uma deformação tridimensional (3D) da coluna vertebral que requer uma avaliação a três dimensões. No entanto, as técnicas convencionais de imagem 3D mostram-se desadequadas e, por isso, as reconstruções são normalmente realizadas recorrendo a radiografias planares, uma no plano frontal e outra no plano lateral. Porém, esta abordagem apresenta desafios. O primeiro prende-se, desde logo, com a extracção de dados 3D a partir de radiografias 2D. Actualmente, este problema é solucionado através de métodos de calibração que recorrem a objectos de calibração radiopacos, e de plataformas rotatórias que visam minimizar os erros de posicionamento do indivíduo. No entanto, estes métodos requerem alterações consideráveis nos protocolos e instalações radiológicas, para além de introduzirem artefactos nas imagens que, por vezes, se sobrepõem a estruturas anatómicas de interesse. Após a calibração das radiografias torna-se necessário resolver um segundo problema, a recuperação da forma da coluna. Os métodos de referência requerem identificação manual de um conjunto extenso de pontos em cada uma das radiografias, o que os torna dispendiosos e propensos a erros. Por este motivo, e pelas razões supramencionadas, os métodos actualmente utilizados para a reconstruação 3D da coluna vertebral não são viáveis para avaliações clínicas de rotina em serviços de radiologia com configurações comuns.

Esta tese assentou na hipótese de que é possível construir modelos geométricos personalizados da coluna vertebral escoliótica a partir de duas radiografias adquiridas em ambiente clínico comum, por meio de métodos computacionais que requerem baixa interacção com o utilizador. Para alcançar este objectivo são propostos dois métodos: um método de calibração geométrica que recorre a uma distância, facilmente medida no local, que visa minimizar o recurso a objectos de calibração; e um método de reconstrução 3D baseado na deformação de um modelo articulado da coluna que visa reconstruções rápidas e exactas.

Mostra-se pela primeira vez, nesta tese, que é possível recuperar a escala da coluna vertebral sem recurso a objectos de calibração, ainda que a introdução de um pequeno objecto permita melhorar a exactidão e robustez das reconstruções 3D. Adicionalmente, mostra-se que o modelo proposto permite melhorar os resultados de calibração de radiografia biplanar baseada na minimização do erro de retro-projecção. O método de calibração proposto requer apenas mudanças pontuais nos protocolos de aquisição radiográfica, introduzindo poucos ou nenhuns artefactos nas radiografias. Conclui-se, assim, que satisfaz os requisitos necessários para utilização em ambientes clínicos comuns.

No que diz respeito à recuperação da estrutura da coluna vertebral, o método proposto alcançou o menor tempo de reconstrução, tendo em conta o tempo de interacção com o utilizador e o tempo de computação da solução. Alcançou ainda, de entre os métodos semi-supervisionados, a maior exactidão na determinação da localização das vértebras, assim como do centro das faces dos corpos vertebrais, tanto em casos de escoliose moderada como severa. Adicionalmente, foi mostrado que o método permite o cálculo de índices clínicos de forma fiável, mesmo quando o input é fornecido por utilizadores não-especialistas. Conclui-se, deste modo, que o método de reconstrução 3D pode ser utilizado para avaliações clínicas de rotina.

Résumé

La scoliose est caractérisée par une déformation par trois dimensions (3D) de la colonne vertébrale qui nécessite une évaluation en trois dimensions. Cependant, les techniques classiques de l'image 3D se révèlent insuffisants, donc les reconstructions 3D sont généralement effectuées à l'aide des radiographies planes, une au plan frontal et une autre dans le plan latéral. Mais cette approche présente des défis. Le premier a trait à l'extraction de données 3D à partir de radiographies 2D. Actuellement, ce problème est résolu par des méthodes de calibrage qui utilisent des objets radio-opaques et des plates-formes rotatives qui visent à minimiser les erreurs de positionnement de l'individu. Cependant, ces méthodes nécessitent de changements considérables dans les protocoles et dans les installations radiologiques, ainsi qu'ils introduisent des artefacts dans les images qui, parfois, chevauchent les structures anatomiques d'intérêt. Après le calibrage des radiographies il est nécessaire de résoudre un deuxième problème, la récupération de la forme de la colonne. Les méthodes de référence exigent l'identification manuelle d'un vaste ensemble de points dans chacune des radiographies, ce qui les rend coûteux et sujets à des erreurs. Pour cette raison, ainsi que d'autres présentées ci-dessus, les méthodes actuellement utilisées pour la reconstruction 3D de la colonne vertébrale ne sont pas réalisables pour les évaluations cliniques de routine dans les services de radiologie qui ont des configurations communs.

Cette thèse était repose sur l'hypothèse que c'est possible de construire des modèles géométriques personnalisés de la colonne vertébrale scoliotique à partir de deux radiographies acquises dans un environnement clinique commun, à travers les méthodes informatiques qui nécessitent une interaction faible avec l'utilisateur. Pour atteindre ce but, deux méthodes sont proposées : une méthode de calibrage géométrique qui utilise une distance, facilement mesurable sur le site, qui cherche à minimiser l'utilisation des objets de calibrage ; et une méthode de reconstruction 3D basée sur la déformation d'un modèle articulé de la colonne qui vise des reconstructions rapides et précises.

Dans cette thèse, il est montré pour la première fois qui c'est possible de récupérer l'échelle de la colonne vertébrale sans l'aide d'objets de calibrage, même si l'introduction d'un petit objet permet d'améliorer la précision et la robustesse des reconstructions 3D. En outre, il est démontré que le modèle proposé permet d'améliorer les résultats de calibrage de la radiographie biplan basée sur la minimisation de l'erreur de rétroprojection. La méthode de calibrage proposée n'exige que des changements minimes dans les protocoles d'acquisition radiographique, en introduisant peu ou pas d'objets dans les radiographies. Il s'ensuit qu'il satisfait les exigences nécessaires pour l'utilisation dans les environnements cliniques communs.

En ce qui concerne la récupération de la structure de la colonne vertébrale, la méthode proposée a atteinte le meilleur temps de reconstruction, en tenant compte du temps d'interaction avec l'utilisateur et le temps de calcul de la solution. A aussi atteinte, parmi les méthodes semi-supervisées, la plus grande exactitude dans l'emplacement des vertèbres et du centre des plateaux vertébraux, à la fois dans le cas d'une scoliose modéré que sévère. En outre, il a été montré que la méthode permet le calcul des indices cliniques de façon fiable, même lorsque l'input est fourni par des utilisateurs non-spécialistes. Par conséquent, on conclue que la méthode de reconstruction 3D peut être utilisée pour les évaluations cliniques de routine.

Contents

1 Introduction **25**

1.1 Motivation . . . 26

1.2 Thesis statement . . . 30

1.3 Thesis outline . . . 31

2 State of the art **33**

2.1 Calibration of spine radiographs . . . 35

2.1.1 Large calibration apparatus . . . 36

2.1.2 Avoiding calibration objects . . . 40

2.1.3 Using small calibration objects . . . 42

2.1.4 Hybrid methods . . . 45

2.1.5 Summary . . . 46

2.2 Construction of geometrical models of the spine from radiographs . 47

2.2.1 Highly-supervised methods . . . 47

2.2.2 Towards unsupervised reconstruction of the spine . . . 55

2.2.3 Semi-supervised spline-based methods . . . 58

2.2.4 Summary . . . 63

3 Geometrical calibration of biplanar radiography **67**

3.1 Goal of the geometrical calibration . . . 67

3.2 The geometry of biplanar radiographs formation . . . 69

3.3 General calibration method . . . 71

3.4 Estimating the calibration parameters using a distance measuring device . . . 73

3.4.1 Pre-calibration procedure . . . 74

3.4.2 Calibration parameters estimation . . . 76

3.5 Algorithms . . . 77

3.5.1 Corners detection . . . 77

3.5.2 Least-squares minimisation . . . 78
3.5.3 Triangulation of point-matches . . . 78
3.6 Method evaluation . . . 79
3.6.1 Materials . . . 79
3.6.2 Pre-calibration procedure . . . 79
3.6.3 Calibration assessment . . . 80
3.7 Results . . . 82
3.7.1 Pre-calibration procedure . . . 82
3.7.2 Errors of the initial guess of the calibration parameters . . . 83
3.7.3 Reconstruction accuracy with ideal point-matches . . . 83
3.7.4 Reconstruction accuracy with non-ideal point-matches . . . 83
3.7.5 Optimisers comparison and effect of constraining focal distance . . . 85
3.8 Discussion . . . 85
3.9 Summary . . . 87

4 Calibration of biplanar radiographs of the spine 89
4.1 Calibration system for 3D reconstruction of the spine . . . 89
4.1.1 Radiographs acquisition procedure . . . 91
4.1.2 Point-matches identification . . . 91
4.1.3 Initial guess of the calibration parameters . . . 92
4.1.4 Scale correction . . . 93
4.2 Simulation studies with *in vivo* 3D data . . . 93
4.2.1 Positioning and size of the calibration object . . . 95
4.2.2 Impact of the calibration components . . . 96
4.2.3 Effect of the precision of the measuring device . . . 97
4.3 *In vitro* validation . . . 97
4.4 Results . . . 98
4.4.1 Positioning and size of the calibration object . . . 98
4.4.2 Method accuracy and impact of the calibration components 100
4.4.3 Effect of the precision of the measuring device . . . 102
4.5 Discussion . . . 103
4.6 Summary . . . 107

5 Statistical-based 3D reconstruction of the spine 109
5.1 Statistical models of the spine . . . 109
5.1.1 Articulated representations of the spine . . . 111
5.1.2 Statistics on articulated models of the spine . . . 112

5.1.3 Deformable articulated model of the spine . . . 113
5.2 Spline guided deformation of the articulated model . . . 115
5.3 Generating plausible spine configurations . . . 116
5.4 Refinement of vertebrae location . . . 118
5.5 Method evaluation . . . 120
5.5.1 Experiments with experimented users . . . 120
5.5.2 Experiments with non-expert users . . . 122
5.6 Results . . . 123
5.6.1 Influence of each term in the cost function . . . 123
5.6.2 User interaction versus reconstruction accuracy . . . 123
5.6.3 Spine reconstruction accuracy . . . 125
5.6.4 Vertebrae position and orientation accuracy . . . 126
5.6.5 Clinical indices accuracy . . . 126
5.6.6 Reconstruction time . . . 127
5.6.7 Results for non-expert users . . . 127
5.7 Discussion . . . 127
5.8 Summary . . . 132

6 Conclusion **133**
6.1 Main contributions . . . 134
6.2 Future work . . . 136

Bibliography **141**

A Deformation of vertebrae models using anatomical landmarks **157**

B Generation of digitally reconstructed radiographs from polygon meshes **159**
B.1 Brief description of the algorithm . . . 159
B.2 Results . . . 161

C 2D/3D registration of vertebrae models on radiographs **163**
C.1 Brief description of the method . . . 163
C.2 Preliminary results . . . 167

List of Figures

1.1 Antero-posterior view of 3D representations of three scoliotic spines . 27
1.2 Construction of a 3D personalised geometrical model of the spine from the two radiographs usually prescribed by physicians. 28
1.3 Illustration of a radiograph of the spine 29

2.1 Clinical indices of the spine: Cobb angle, kyphosis, lordosis, spinal length and height 34
2.2 Clinical indices of the plane of maximum deformity 34
2.3 Illustration of the six anatomical points that are typically used to represent a vertebra 35
2.4 Large calibration apparatus proposed by Dansereau and Stokes (1988) 37
2.5 Representation of the calibration apparatus proposed by Dansereau and Stokes (1988) 37
2.6 Large calibration apparatus proposed by Dumas et al. (2003) 39
2.7 The EOS™ system from BioSpace Med, France 40
2.8 Setup of the radiological environment for surgery when using the small calibration object proposed by Cheriet et al. (1999b) 42
2.9 Small calibration object proposed by Cheriet et al. (1999b) 42
2.10 Small calibration object proposed by Kadoury et al. (2007b) 43
2.11 Radiographs showing the calibration object proposed by Kadoury et al. (2007b) 44
2.12 Calibration vest proposed by Cheriet et al. (2007) 45
2.13 Stereo-corresponding anatomical landmarks 48
2.14 Generic model of a lumbar vertebra 48
2.15 Non stereo-corresponding anatomical landmarks 49
2.16 Refining the model of a generic lumbar vertebra 51
2.17 Illustration of the reconstruction method proposed by Pomero et al. (2004) 53
2.18 Illustration of the deformable model used in Benameur et al. (2003) . 56

2.19 Deformable model of the spine used in Benameur et al. (2005) 57
2.20 Classification of curvature type and severity taken into account by the method proposed by Vaiton et al. (2004) 59

3.1 Illustration of a biplanar radiographic exam on a convectional system 68
3.2 Coordinate system of reference and representation of the intrinsic parameters 70
3.3 Illustration of a conventional radiographic imaging system with a rangefinder attached 74
3.4 Radiographs of the phantom for the pre-calibration procedure 75
3.5 Relation between d_d and d_s for $d_m = 909$ mm on a real imaging system 76
3.6 Interception of the relations between d_d and d_s for eight radiographs of the phantom acquired at different distances d_m on a real imaging system 77
3.7 Example of three radiographs of the phantom and the initial guess for the orientation and translation parameters 80
3.8 Histogram of 3D reconstruction errors 84
3.9 3D reconstruction error vs. noise in point-matches location 84
3.10 Comparison of optimisers and effect of the range of the bounds for the focal distance (f) on the trust-region optimiser 85
3.11 Reconstruction of a case where the Levenberg-Marquardt algorithm failed to capture the scene geometry 86
3.12 Illustration of the effect of the range of the bounds for the focal distance (f) when using the trust-region optimiser 86

4.1 Data flow diagram of the 3D reconstruction system 90
4.2 Illustration of a conventional radiographic imaging system with a rangefinder attached 92
4.3 Illustration of a simulation with an *in vivo* CT of the spine 94
4.4 Radiographs and CT reconstruction of the dried vertebrae 98
4.5 Effect of the orientation of the calibration object on the 3D reconstruction error 99
4.6 Illustration of the effect of the position of the calibration object on the 3D reconstruction error 99
4.7 Effect of the size of the calibration object on the 3D reconstruction error 100
4.8 Simulation results for 3D reconstruction accuracy 101
4.9 *In vitro* results of the proposed calibration method 101

4.10 Effect of the precision of the distance measuring device on the 3D reconstruction of the spine . . . 103

5.1 The spine as an articulated model . . . 111
5.2 Principal deformation modes of the statistical model of the spine . . . 113
5.3 Graphical user interface (GUI) designed for identifying the splines . . . 114
5.4 Illustration of the fitting error . . . 117
5.5 Using the location of control points to improve fitting . . . 119
5.6 Impact on the 3D reconstruction accuracy of the terms of the cost function . . . 124
5.7 Mean RMS 3D reconstruction error vs number of control points per spline and per radiograph . . . 124
5.8 Comparison of 3D reconstructions using the proposed method vs reference method . . . 125
5.9 Projection of the 3D reconstructions onto the original radiographs . . . 126

6.1 Gradient of a posterio-anterior radiograph of a pedicle . . . 137
6.2 Preliminary results of 2D/3D registration . . . 138

A.1 Generic 3D model of a vertebra . . . 158
A.2 Illustration of a spine reconstruction where vertebrae were rendered using polygon meshes. . . . 158

B.1 Formation of a Digitally Reconstructed Radiograph (DRR) from a polygon mesh . . . 160
B.2 Digitally reconstructed radiographs (DRR) obtained by the proposed algorithm . . . 162
B.3 DRR of a complete spine and comparison with the original radiograph 162

C.1 Response of the selected metric (GDG) to 2D translation on the postero-anterior view . . . 165
C.2 Response of the selected metric (GDG) to 2D translation on the lateral view . . . 165
C.3 Preliminary results of 2D/3D registration . . . 166

List of Tables

2.1 Comparison of *in vitro* 3D reconstruction errors between highly-supervised methods . . . 51
2.2 Comparison of *in vitro* orientation errors between highly-supervised reconstruction methods . . . 52
2.3 Comparison of *in vivo* 3D reconstruction errors between highly-supervised methods . . . 54
2.4 Orientation errors for SA when using NSCP reconstructions as basis . 54
2.5 Synthesis of studies evaluating the accuracy to recover the global shape of the spine . . . 64
2.6 Synthesis of studies evaluating the accuracy of vertebrae shape reconstruction . . . 65

3.1 Translation and orientation absolute errors of the initial guess . . . 83
3.2 3D reconstruction error vs. number of point-matches and noise in point-matches location . . . 84

4.1 Parameters' noise used on the simulations . . . 95
4.2 Relation between the calibration object position and the subject's orientation . . . 99
4.3 Simulations results for absolute orientation errors . . . 101
4.4 Comparison of errors *in vitro* with related work . . . 102

5.1 Reconstruction errors by an experimented operator . . . 128
5.2 Location and orientation errors by an experimented operator . . . 128
5.3 Evaluation results of clinical indices by an experimented operator . . . 128
5.4 Reconstruction time comparison with other spline-based methods . . . 128
5.5 Reconstruction errors by non-expert users . . . 129
5.6 Evaluation results of clinical indices by non-expert operators . . . 129

B.1 Peformance of the DRR generator implemented in OpenGL GLSL . . 162

Chapter 1

Introduction

Nowadays, three-dimensional (3D) medical imaging techniques, such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI), are easily available for assessing both bone and soft tissue. However, 3D imaging techniques usually require subjects to be lying down. If for several organs this is not a problem, the same may not be said regarding the spine. The spine is an articulated structure that changes its global shape depending of the subject position. Deformities of the spine, such as scoliosis, are evaluated in standing position, rendering conventional 3D imaging techniques unsuitable for this task. Currently, capturing the spine in standing position with adequate resolution is only viable using planar radiography and, thus, this remains the standard imaging technique for clinically assessing scoliosis (Cassar-Pullicino and Eisenstein, 2002; Fairbank, 2002; Greenspan, 2004). Still, proper evaluation of scoliosis requires 3D metrics of the spine that cannot be assessed directly on planar radiographs (Stokes, 1994). Therefore, several methods have been proposed to reconstruct the spine in 3D from two or more planar radiographs. In this thesis we will first show that the current methods are not viable for standard clinical environments, since they require several changes to radiological setups and/or require high levels of user-interaction by an expert. We will then propose computer methods that aim at making possible to have 3D reconstructions of the spine from two radiographs (biplanar radiography) on standard clinical setups, with reconstruction times compatible with the bustle of routine clinic, while remaining sufficiently

accurate for properly evaluating scoliosis.

The remainder of this chapter starts with a more detailed explanation of the motivation for this thesis, followed by the thesis statement and a brief description of the remaining chapters.

1.1 Motivation

According to the Scoliosis Research Society, scoliosis is defined as a lateral deviation of the normal vertical line of the spine which, when measured by radiograph, is greater than ten degrees (figure 1.1) (Working Group on 3-D Classification, 2000). In fact, this traditional definition is limited, since the deformity occurs in varying degrees in all three anatomical planes (Stokes, 1994). Additionally, vertebrae are rotated about the spine midline (axial rotation), and other structures are also affected, such as the rib cage.

It is estimated that scoliosis affects 1% to 2% of the adolescents (Velezis et al., 2002; Karachalios et al., 1999). When left untreated, this medical condition can lead to progression of the deformity, back pain, cardiopulmonary problems, and psychosocial concerns (Weinstein et al., 2008). Treatment recommendations include watchful waiting, physiotherapy, bracing, and surgery, with the later being only indicated for severe cases (Weinstein et al., 2008). The most used non-operative therapy is bracing (Rowe et al., 1997; Labelle et al., 2007). It is utilised during adolescence and aims at preventing progression of the deformation until the subject reaches skeletal maturity, at which time the risk of curve progression, and hence the risk of surgery, greatly diminishes (Weinstein et al., 2008). Surgery is also advised before adulthood due to more flexibility of the spine and decreased rate of complications. Thus, given the small time window for effective therapy, an accurate assessment of the deformity is paramount to fully characterise the pathology and to provide the most appropriate treatment for each patient.

For assessing three-dimensional deformations of the spine like scoliosis, one would expect that 3D examinations, such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI), would be the elected ones. However, these methods have several disadvantages that make them unsuitable to recover the global spine geometry. First, both methods require the subject to be laying down, which alters the spine posture (Yazici et al., 2001). When evaluating scoliosis the subject should be in erect position (Cassar-Pullicino and Eisenstein,

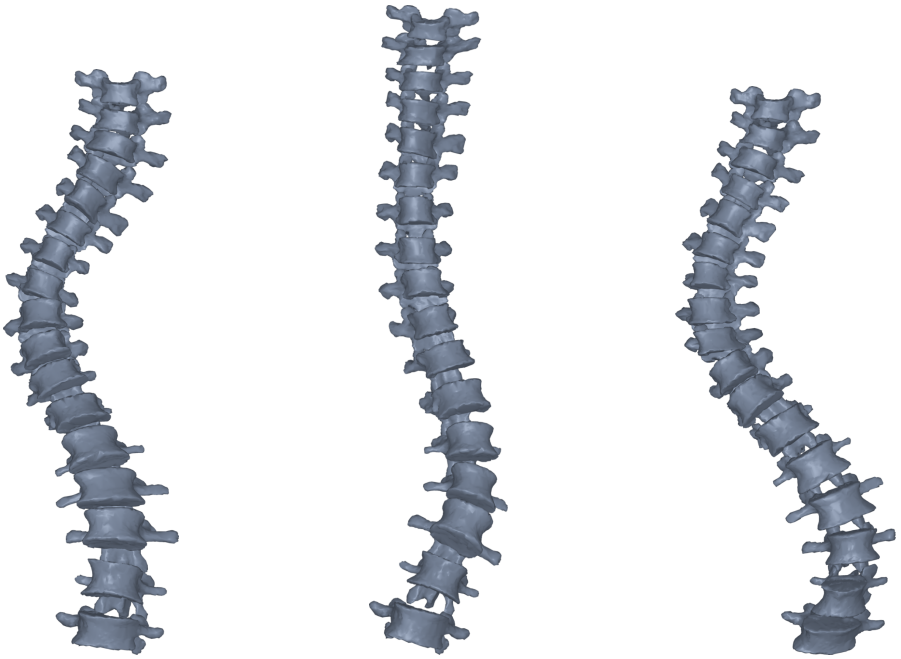


FIGURE 1.1: Frontal (antero-posterior) view of 3D representations of three scoliotic spines. The pathological curve on the frontal plane is visible on all of them. A non-pathological spine should present no curvature on this view.

2002; Fairbank, 2002; Greenspan, 2004). Additionally, both methods are expensive, especially MRI. CT scans are sensitive to bone tissue (Rogers, 1998), but it is impracticable to use CT for scoliosis evaluation due to the high doses of radiation needed to cover the entire spine (Levy et al., 1996). CT scans are mainly used to exam a limited set of vertebrae when highly detailed reconstructions are required (e.g. vertebrae fracture detection) (Greenspan, 2004; Rogers, 1998). Unlike CT, MRI is a non-invasive technique that is especially adequate for analysing soft tissue, bone marrow, cartilage and ligaments (Greenspan, 2004). Unfortunately, this powerful imaging technique is typically unable to capture cortical bone (the outer layer of bone that defines its surface) (Rogers, 1998) and therefore it is not adequate for capturing the geometry of bone structures with detail. Nevertheless, MRI is an important tool for assessing neurological complications that may be originated by the irregular curvature of the spine, and is often used as a complementary diagnosis examination (Cassar-Pullicino and Eisenstein, 2002; Greenspan, 2004; Fairbank, 2002).

For the above reasons, planar radiography is still the gold-standard technique for evaluating scoliosis. Planar radiography is much less invasive than CT while remaining sensitive to bone structures, which allows to acquire the entire spine.

The standard radiographs that are prescribed are one lateral radiograph (left-right or right-left) and one frontal (anterior-posterior (AP) or posterio-anterior (PA)) (Cassar-Pullicino and Eisenstein, 2002; Fairbank, 2002; Greenspan, 2004). These radiographs enable to evaluate the spine curvature in two orthogonal planes but still do not provide 3D information about the pathology. Several clinical indices, such as the orientation of the maximum plane of deformation (Stokes, 1994), are not possible to assess. According to The Scoliosis Research Society Working Group on 3D Terminology of Spinal Deformity "Two-dimensional measurements always simplify the real deformity of the spine" (Stokes, 1994).

For dealing with these issues, several methods have been proposed for reconstructing the spine in three dimensions from the two radiographs that are usually prescribed (figure 1.2). These methods try to give physicians a 3D visualisation and 3D measurements without subjecting the patient to further examinations and, consequently, to further radiation. Three-dimensional reconstructions from planar radiographs have allowed to properly evaluate the effect of different therapeutic approaches (e.g. Boston braces (Labelle et al., 1996) and surgery (Labelle et al., 1995b)), predict the progression of scoliosis (Villemure et al., 2001) and design more effective braces (Labelle et al., 2007).

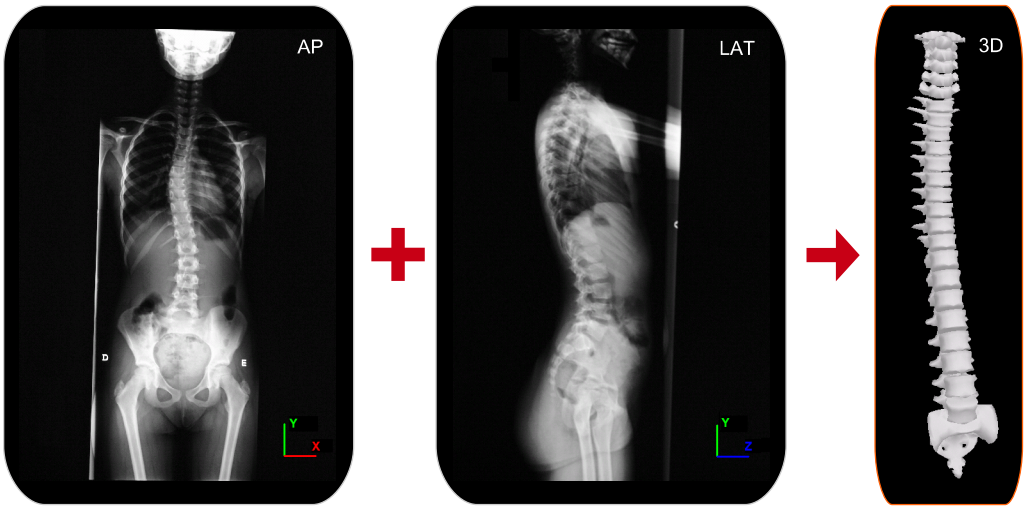


Figure 1.2: Construction of a 3D personalised geometrical model of the spine from the two radiographs usually prescribed by physicians.

One of the main challenges when extracting geometrical information from radiographs is mapping radiographs' coordinates into real-world coordinates. Planar radiographs are 2D projections with no depth information. A static x-ray source emits diffuse radiation for capturing the target area with a single shot.

The radiation passes through the subject body and is then captured by a x-ray detector (figure 1.3). Unfortunately, this method causes non-uniform changes in the size of the radiographed objects that depend of the objects' position relatively to the position of the x-ray source and to the position of the x-ray detector. In spine radiographs the scaling effect has a great impact because the target area is very large and because the distance between vertebrae and the detector is not constant due to the curvatures of the spine, either natural or pathological (figure 1.3). Additionally, the position of the x-ray source should be adjusted according to the subject's anatomy (e.g. infant vs. adult), and subject's position and orientation on the radiographs is not accurately known, making difficult to obtain accurate 3D reconstructions.

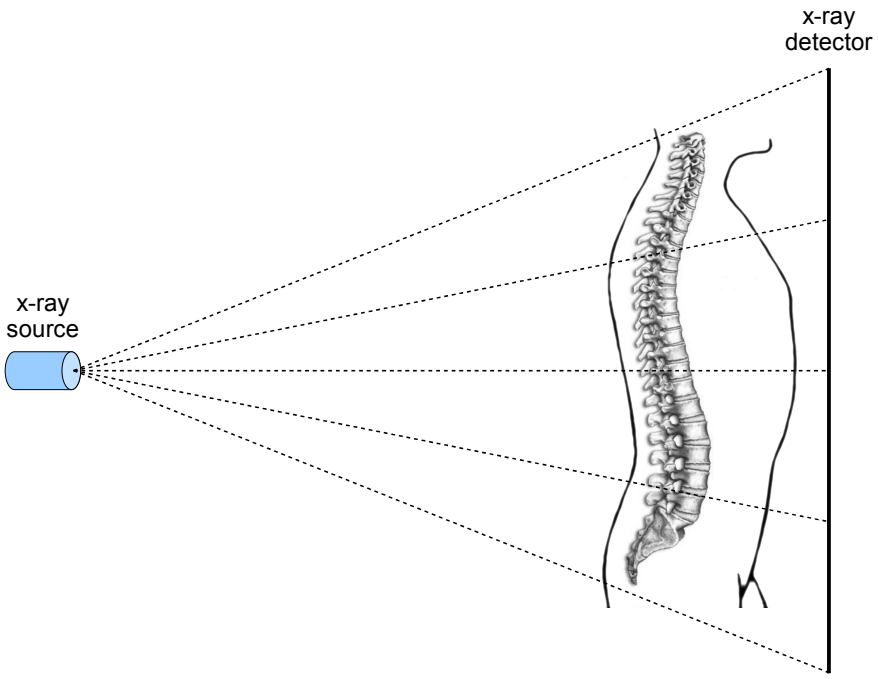


FIGURE 1.3: Illustration of a radiograph of the spine. Spine radiographs cover a very large area, and distances between vertebrae and the detector are not constant, which results in scale distortions caused by the 3D-2D projection.

For tackling the aforementioned problems, calibration methods (detailed in chapter 2) were developed for spine radiography that model the projection of 3D real world coordinates into 2D radiograph coordinates. For determining the geometrical parameters of the model, typically, a set of calibration objects are fixed to the environment and/or to the subject. Unfortunately, the most widely used calibration apparatus require considerable adaptations to the radiological environment that are neither practical nor affordable. In particular, calibration

apparatus have to be installed in the radiological environment that are built big enough to surround the subject (Dansereau and Stokes, 1988; Dumas et al., 2003), or a rotatory platform with calibration objects is required (Cheriet et al., 2007). These solutions require dedicated equipment that makes them unsuitable for standard clinical environments, and introduce several artifacts in the radiographs. For overcoming these limitations, attempts have been made of using small calibration objects (Cheriet et al., 1999b; Kadoury et al., 2007b) but they require extensive user interaction making them unsuitable for routine clinical exams. Finally, other approaches try not to use any calibration object at all, but results shows that scale cannot be recovered under these conditions (Cheriet et al., 1999a; Kadoury et al., 2007a).

Concerning the 3D reconstruction of the spine, the gold-standard methods are based on the manual identification of several anatomical landmarks on each vertebrae for each radiograph (Aubin et al., 1995; Mitton et al., 2000). Consequently, 3D reconstructing the spine of a subject with these methods can take more than an hour, making them unsuitable for routine clinical environment. Attempts of eliminating supervision were achieved but only for the lower part of the spine (Benameur et al., 2005), which is not sufficient for clinical assessment. Therefore, efforts in the last years were directed to semi-supervised approaches in order to have reliable input for guiding reconstructions, while not requiring exhaustive user-interaction. Nevertheless, most of the approaches still require considerable user-interaction (Pomero et al., 2004; Dumas et al., 2008; Humbert et al., 2009) or considerable computation time (Kadoury et al., 2009a) and their accuracy determining clinical indices was not proven.

1.2 Thesis statement

This thesis states that it is possible to build personalised geometrical models of scoliotic spines from two radiographs acquired in standard clinical environments by means of computational methods that require limited user input.

Concerning calibration, our efforts are directed to minimise adaptations to standard clinical setups and, therefore, the use of calibration objects is avoided or minimised. Ideally, scale should be recovered without calibration objects, something that was not previously reported in the literature. Recent work has shown that measuring several angular clinical indices is possible without calibration objects (Kadoury et al., 2007a) and that user interaction can be substantially

decreased using image processing techniques (Kadoury et al., 2010); however, accuracy of 3D reconstructions remains noticeable inferior to the remaining methods.

Regarding the 3D reconstruction of the spine, the main goals are achieving limited user-interaction with low total reconstruction time, while delivering both accurate reconstructions as well as accurate clinical indices for both moderate and severe cases of scoliosis. Achieving such goals are required for routine clinical assessment and were also not reported in the literature. Additionally, decreasing the level of expertise that is usually required by reconstruction methods would facilitate their introduction in clinical institutions.

1.3 Thesis outline

The reminder of this thesis is outlined as follows.

Chapter 2 introduces scoliosis evaluation and reviews the main literature concerning both geometrical calibration of multi-planar radiography of the spine as well as 3D reconstruction methods.

Chapter 3 proposes a novel calibration method for general biplanar radiography that extends a well known calibration technique used on spine biplanar radiography by introducing a distance in the model that can be easily measured on site. This chapter includes details concerning the implementation of the method in a radiologic system as well as experiments with a phantom of known geometry on a real environment.

Chapter 4 addresses the problem of adapting the method proposed in the previous chapter for the case of spine radiography in clinical environment. The behaviour of the method under these conditions was studied using computer simulations with *in vivo* 3D data of the spine. The method is finally validated *in vitro* showing that the proposed extension enables to recover the scale of reconstructions without requiring calibration objects.

Chapter 5 proposes a novel method for spine 3D reconstruction based on the deformation of an articulated model of the spine towards fitting limited user input. The method is validated *in vivo* and results are thoroughly compared with related work showing that it delivers the fastest reconstructions while providing accurate clinical indices, even by non-expert users.

Finally, chapter 6 ends this thesis by presenting general conclusions and pointing out future research directions.

Chapter 2

State of the art

Three-dimensional reconstructions from multi-planar radiographs allow to recover the 3D structure of the spine with subjects on standing position, which is required for properly evaluating scoliosis (Stokes, 1994; Yazici et al., 2001). Ideally, the 3D reconstruction would target to recover vertebrae's shape. However, reconstructing six 3D points per vertebra (Labelle et al., 1995a; Delorme et al., 2003) already allows to estimate the clinical indices suggested by the Scoliosis Research Society (Stokes, 1994). Some of these indices are angular metrics (e.g. Cobb angle, kyphosis and lordosis on figure 2.1, and orientation of the plane of maximum deformation on figure 2.2), while others are based on Euclidian distances (e.g. spinal length and height illustrated on figure 2.1). The six anatomical points that are commonly used to represent a vertebra are the centre of the superior and inferior endplates, and the superior and inferior extremities of the pedicles (figure 2.3). The endplates' landmarks enable to determine the location of the vertebral body and, consequently, the vertebral body line (Stokes, 1994), also known as spine midline. The vertebral body line is the curved line that passes by all the vertebral bodies, and it is used to calculate most of the clinical indices of the spine (Stokes, 1994). The pedicles, together with the endplates, allow to determine vertebrae pose (Stokes, 1994), especially axial rotation that cannot be calculated from the vertebral body line. Additional anatomical points may help recovering vertebrae's shape (e.g. Mitton et al. (2000); Delorme et al. (2003)) and evaluate local deformations of vertebrae (e.g. Aubin et al. (1998)).

Source: adapted from Bellefleur (2001)

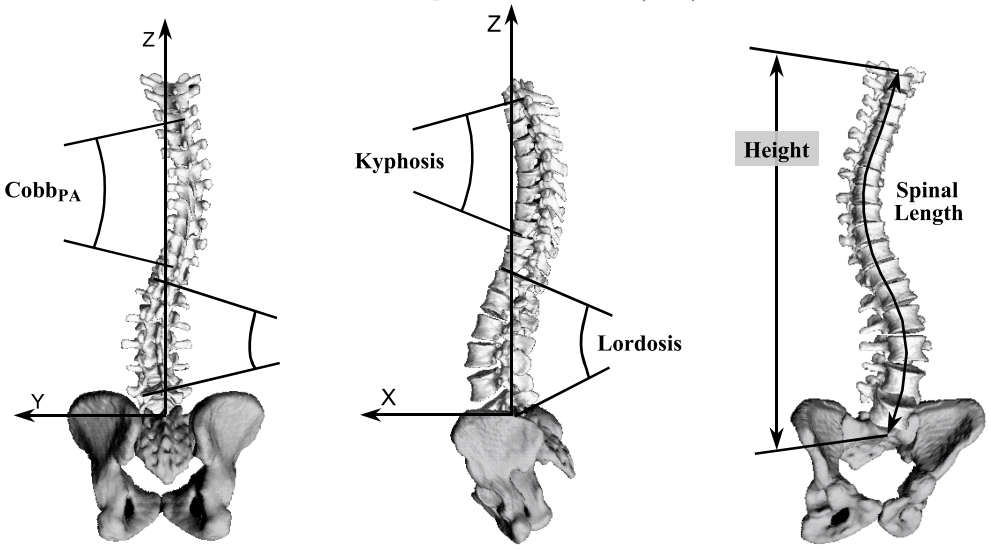


Figure 2.1: Illustration of clinical indices of the spine: maximum Cobb angle on the posterio-anterior plane (Cobb_{PA}), kyphosis and lordosis (lateral plane), and 3D spinal length and height.

Source: adapted from Bellefleur (2001)

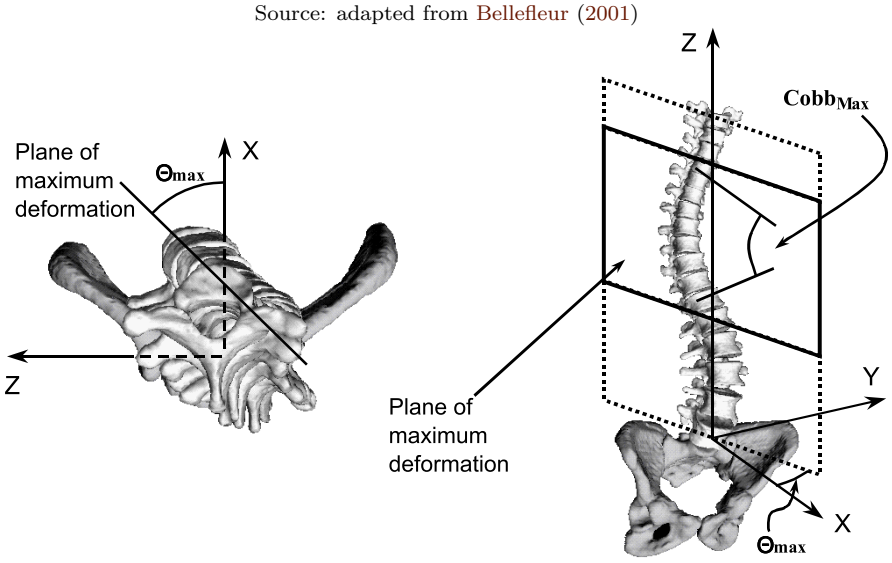


Figure 2.2: Illustration of the plane of maximum deformity, its orientation (Θ_{max}) and the Coob angle on this plane (Cobb_{Max}) – top view (on the left), and perspective view (on the right).

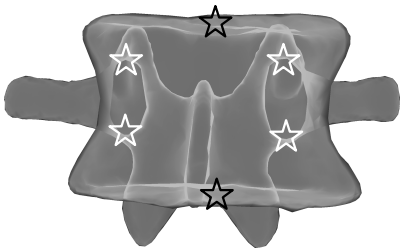


FIGURE 2.3: Illustration of the six anatomical points that are typically used to represent a vertebra (frontal view). (White stars represent pedicles' landmarks while black stars show the centres of the endplates. The model of the vertebra is semi-transparent for illustration purposes.)

Performing 3D reconstructions from radiographs requires two steps: (i) geometrical calibration and (ii) recovering the shape of the spine. The first step estimates the geometrical parameters of the system, which enable to calculate 3D data from 2D data. The second step captures the structure of the spine, e.g. 6 points per vertebra, or even a 3D model of the shape of each vertebrae.

In the reminder of this chapter the main methods for both geometrical calibration of multi-planar radiography of the spine and shape recovery of the spine are reviewed and compared. The main advantages and disadvantages of each approach are highlighted as well as limitations of the current methods concerning their utilisation on standard radiological setups for routine clinical assessment. During the reading of this chapter, the reader is invited to consult tables 2.5 and 2.6 (on pages 64–65) that summarise the main methods reviewed here.

2.1 Calibration of spine radiographs

Radiography calibration is a very similar problem to camera calibration from computer vision. The goal is to find the calibration matrix that maps world coordinates into image coordinates. The calibration matrix is composed by two set of parameters: intrinsic and extrinsic. Intrinsic parameters are related to the properties of the system, and extrinsic parameters are related to the position of the target in relation to the camera (or, in this case, the x-ray system). By determining these parameters one is able to calculate the 3D coordinates of a point that is visible in two or more planar projections, as well as to project a 3D point to the images' planes.

Several techniques from the computer vision community may be used to try

solving the radiography calibration problem with few adjustments. These techniques typically require the matching between features (e.g. points) across different views. Such features may be provided either by the natural content of the images or by introducing easily identified objects (calibration objects) on the scenes. The classic approach for spine radiography involves a large calibration object with sufficient dimensions to surround a person. This approach yields accurate results but has several inconveniences, such as costs and portability, among others. For dealing with these problems, in the last years there have been efforts for reducing the size of calibration objects or even eliminate them at all. In the next subsections we will review the most common approaches and the ones that have emerged in the last years.

2.1.1 Large calibration apparatus

One of the most employed approaches in the literature for calibrating multiplanar radiographs when constructing geometrical models of the spine is using a calibration object that is large enough to surround the subject's body, e.g. Aubin et al. (1997); Mitton et al. (2000); Mitulescu et al. (2001, 2002); Dumas et al. (2003); Benameur et al. (2003, 2005); Pomero et al. (2004). The first method which fits in this class was proposed by Dansereau and Stokes (1988). The authors proposed using a very well known technique in computer vision: Direct Linear Transformation (DLT) (Abdel-Aziz and Karara, 1971). Basically, the DLT algorithm finds the intrinsic and extrinsic parameters by establishing a transformation that translates the world coordinates of known control points into image coordinates (Wohler, 2009). For accomplishing this, it requires a minium of six points (if lens distortion parameters are not required) with known 3D location in the world and known 2D location in the image. Therefore, when applying DLT to radiography it is necessary to have a calibration object composed by a set of radiopaque pellets that should be fixed to the radiological environment in order to know their exact 3D location. In practice, more than six pellets are used for compensating identification errors on radiographs, which adds robustness to the algorithm. Unfortunately, DLT presents significant extrapolation errors (Wood and Marshall, 1986). Thus, for accurately estimating 3D coordinates, Dansereau and Stokes proposed a calibration object that was built large enough to contain, within its limits, the anatomical structures that are being reconstructed (Dansereau and Stokes, 1988). In figure 2.4 it is possible to observe the actual design of the calibration apparatus. This apparatus includes a rotatory platform

for minimising patient movements between radiograph acquisitions, and for keeping an approximately constant distance between the x-ray source and the spine (figure 2.5).

Source: adapted from Cheriet et al. (1999a)



FIGURE 2.4: Large calibration apparatus proposed by Dansereau and Stokes (1988).

Source: adapted from Cheriet et al. (1999a)

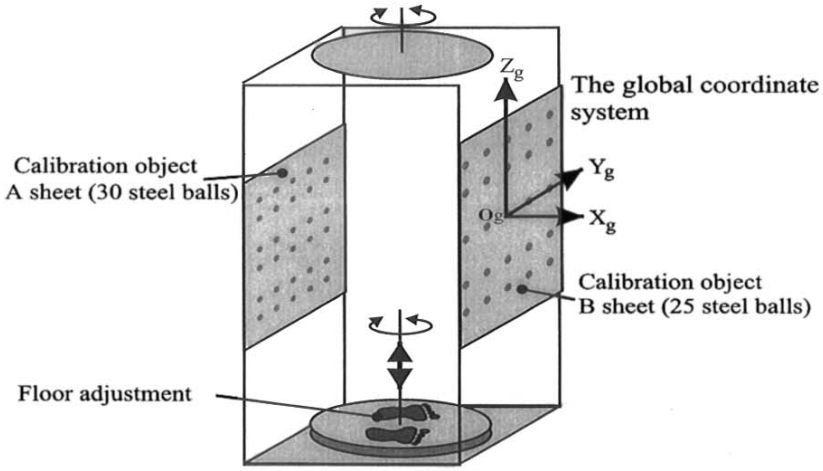


FIGURE 2.5: Representation of the calibration apparatus proposed by Dansereau and Stokes (1988).

In an *in vitro* study (Aubin et al., 1997) of the method proposed by Dansereau and Stokes (1988), 3D triangulation errors of manually identified landmarks (21 per vertebra) were of 2.1±1.5mm (mean ± S.D.). For *in vivo* studies, the best results were achieved by Mitulescu et al. (2002) that measured an error of 1.5mm

mean point to surface distance when compared to CT scans reconstructions with accuracy of 1.1mm. When comparing *in vivo* reconstructions between multiplanar radiography and CT or MRI, each vertebra has to be rigidly aligned before computing reconstruction errors because their position and orientation is different. Therefore, *in vivo* studies only allow to validate vertebrae shape rather than the global shape of the spine.

Despite of the good results achieved with this calibration technique, it has several drawbacks. For starting, it requires considerable changes in standard radiographic environments that may not be affordable or even possible to implement. Additionally, the calibration pellets overlap anatomical structures on the radiographs. Besides these problems, Cheriet et al. (1999a) documented that the physical characteristics of the calibration object make some subjects to feel fear or uncomfortable, requires some patients to kneel, and are not adequate for patients lying down (e.g. during surgery) and neither for patients that are not able to stand up (e.g. patients in wheelchairs). Moreover, since the pellets are fixed to the environment (and not the subject), the method is not able to coupe with patient motion between radiographs.

For trying to tackle some of the above problems, Dumas et al. (2003) developed a new calibration device and a calibration method that was not based on DLT. Instead, the authors developed a simplified geometrical model of the radiological environment and made some assumptions that enabled to get a set of calibration equations. The main assumptions are: (i) the image reference plane must be parallel to the global reference frame in both radiographs (frontal and lateral), and (ii) the x-ray source must remain in the same position between radiographs. Having these assumptions in mind and the proposed calibration equations, the authors developed the calibration device that is presented in figure 2.6. As one may see, now the calibration pellets rotate with the subject, who is no longer fully enclosed by the calibration apparatus. This system is therefore more patient-friendly and has the additional advantage of decreasing the number of pellets that superimpose bone structures. For evaluating the system accuracy, this method was compared with the previous one. In particular, the authors used the same dried vertebrae and the same reconstruction algorithm presented in Mitulescu et al. (2001), where calibration was performed with the DLT-based method. The results were very similar: the new system was able to achieve a mean point to surface distance of 1.2mm (against 1.1mm), RMS of 1.6 (against 1.4mm) and maximum error of 6.4mm (against 7.8mm).

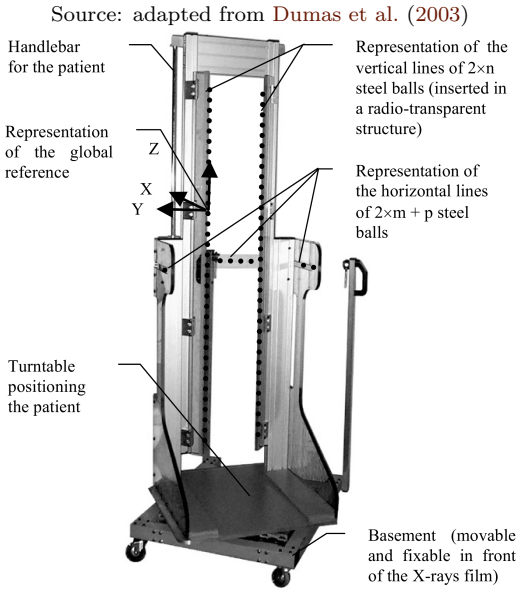


FIGURE 2.6: Large calibration apparatus proposed by Dumas et al. (2003). (n=32, m=8, p=11).

Recently, a new commercial imaging system for biplanar radiography, EOS™ from BioSpace Med, France (figure 2.7), was introduced in the market (Dubousset et al., 2005). This system provides a new approach where the two radiographs are acquired simultaneously, which enables to know the exact relative rotation between the two views and eliminates the problem of patient motion between radiographs. Additionally, instead of acquiring the image with a single shot like in conventional radiography, the system is capable of sweeping the complete body from head to toes, which eliminates distortion in the vertical axis. However, this sweeping takes several seconds and, thus, introduces a new problem that is not present on conventional radiography: patient motion during the acquisition. The system is calibrated initially after installation and then calibration parameters remain the same for all patients. Extrinsic parameters are roughly guaranteed by feet markers on the floor, which may be inaccurate. Despite there is no vertical distortion on radiographs, there is a considerable distortion on the horizontal axis, which depends of the position of the anatomical structures with relation to the x-ray sources and detectors. Unfortunately, there are no accuracy studies concerning the global location of vertebrae's landmarks. Only accuracy of vertebrae's shape was evaluated after rigidly aligning each vertebra with CT (Bras et al., 2003). Results on this topic were slightly better than results of the previous methods, but there are some issues that make this comparison difficult: the reconstruction method proposed in Bras et al. (2003) has non-negligible differences

to the method used for evaluating shape accuracy on the previous calibration systems, the set of vertebrae is not the same, and the validation technique is less accurate. Nevertheless, the errors of point to surface distance were, mean 0.9mm (against 1.1mm), RMS 1.2mm (against 1.4mm), and maximum error 5.8 (against 6.4mm). In logistical terms, the main disadvantages of this system are its costs and the need to have additional space to install it, since it does not replace standard x-ray imaging systems.

Source: adapted from Valenza (2009)



Figure 2.7: The EOS™ system from BioSpace Med, France. This system is able of capturing frontal and lateral radiographs of the patient simultaneously.

2.1.2 Avoiding calibration objects

As we have seen in the previous section, calibration objects have several disadvantages, especially if they have large dimensions. In order to avoid using calibration apparatus, Cheriet et al. experimented calibrating radiographs without using any apparatus. The input of the proposed method was (i) a set of stereo-corresponding points (points that are visible on both images) that were manually identified over specific regions of every vertebra, and (ii) an initial guess of the calibration parameters. These parameters were then optimised in order to minimise the mean square distance between the observed and analytical projections of the marked points. The authors experimented this method with angiography (Cheriet and Meunier, 1999) and spine radiography (Cheriet et al., 1999a). When calibrating angiography images, the authors had access to a very good estimation of some of the calibration parameters because the angiography machine gives information about the x-ray source position and the angle between the two

acquisitions. This good estimation enabled to reduce the search space enough to find accurate calibration parameters (Cheriet and Meunier, 1999). However, when applying the same method to spine radiography, the authors did not have access to the x-ray source position, and neither to the geometric transformations between the two acquisitions. For filling this gap, the authors collected the calibration parameters of several radiographs taken with a large calibration object (Dansereau and Stokes, 1988) in order to have an average solution and a standard deviation. Using this information to choose an initial guess for the optimisation algorithm and to narrow the search space was proved to be insufficient for finding a good solution of the calibration parameters (Cheriet et al., 1999a). Having failed to calibrate spine radiographs, the authors attempt to find out the minimum absolute information that a calibration object should provide. For this task they used the calibration object presented in the previous section proposed by Dansereau and Stokes (1988). The conclusion was that for accurate results the algorithm should have access to the absolute coordinates of three coplanar points and to the absolute measure of four distances.

Using the same method, Kadoury et al. (2007a) tried to perform geometric spine reconstructions without using any calibration object. The results were that angular measures of the spine were recovered with small errors, but the same did not happen to metrics based on Euclidian distances, such as the spinal length, which had an error of 14.19 ± 8.00mm. This shows that the method is not able of handling the scaling effect, because scaled objects present the same angular measures but their Euclidian dimensions are obviously different. The method was validated using a rotatory platform for patient positioning that minimises patient motion and change of posture between acquisitions, which does not enable to generalise results for standard clinical setups.

Not using calibration objects and not using rotatory platforms has several advantages, such as low costs, minimum adaptation of the radiological environment, minimum constrains for the patient position (e.g. the patient may be seated), and radiography images remain unaltered. To the best of our knowledge, no method as proved to be accurate in calibrating spine radiographs without using calibration objects and without a rotatory platform. Moreover, the methods described here require considerable user intervention.

2.1.3 Using small calibration objects

Having failed to develop a method that does not rely on calibration objects, a small calibration object was constructed (figures 2.8 and 2.9) in order to obtain accurate calibrations during surgery Cheriet et al. (1999b). The calibration algorithm was the same suggested in Cheriet et al. (1999a) and according to Novosad et al. (2004) the RMS reconstruction error of the spine using this calibration method is 2.5mm. This error is higher than the RMS reconstruction errors achieved using large calibration apparatus, which justifies the preference for large calibration objects when they are available. However, this error is still admissible for several clinical applications, especially when it is not possible to have patients standing up, like in surgery.

Source: adapted from Novosad et al. (2004)

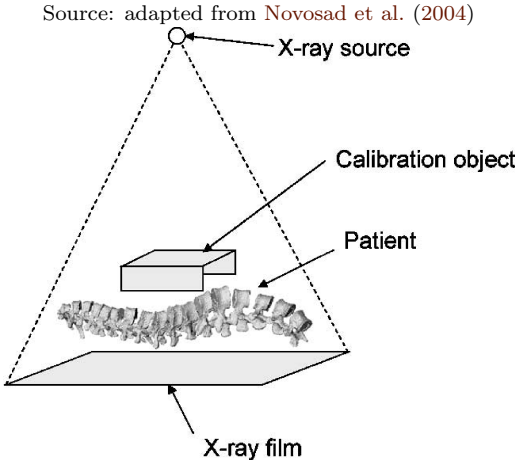


Figure 2.8: Setup of the radiological environment for surgery when using the small calibration object proposed by Cheriet et al. (1999b).

Source: adapted from Novosad et al. (2004)

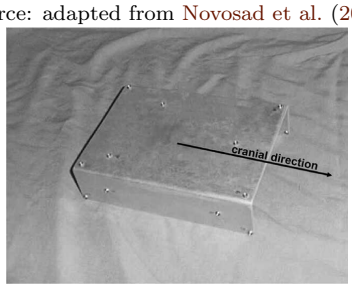


Figure 2.9: Small calibration object proposed by Cheriet et al. (1999b) with a set of 15 steel pellets of known position.

Kadoury et al. (2007b) also developed a small calibration object, but now with the intent of being portable. The object has a set of four pellets with known posi-

tions, and it is placed in a vest worn by the patient during both acquisitions (figures 2.10 and 2.11). The proposed method assumes the weak-perspective camera model (Hartley and Zisserman, 2003) to extrapolate the geometrical parameters of the radiological environment from the location of the calibration pellets on the two images. This technique does not provide accurate calibrations but enables to have a reasonable guess of the parameters values. In order to improve calibration, the authors use the optimisation algorithm proposed by (Cheriet et al., 1999a), which needs a set of stereo-correspondent points. In this work, the authors used a method for constructing the geometrical model of the spine that needs a set of 6 points per vertebra to be manually identified on both images. The authors take advantage of these points and use a subset of them to feed the optimisation algorithm. Using this approach the authors were able to improve the results of their previous work (Kadoury et al., 2007a) where they did not use any calibration object. In particular, they were able to reduce the error when calculating metrics based on Euclidian distances, such as the spinal length that previously scored 14.19 ± 8.00mm and now scored 2.05 ± 1.03mm. This indicates that this method is couping with the scaling effect problem. However, angular clinical indices obtained inferior results (although acceptable). Unfortunately, the authors did not publish any data related to the accuracy of spine reconstruction and therefore a comparison with the other methods on this subject is not possible. Again, the method was validated using a rotatory platform for proper patient positioning, which does not enable to generalise results for standard clinical setups.

Source: adapted from Kadoury et al. (2007b)

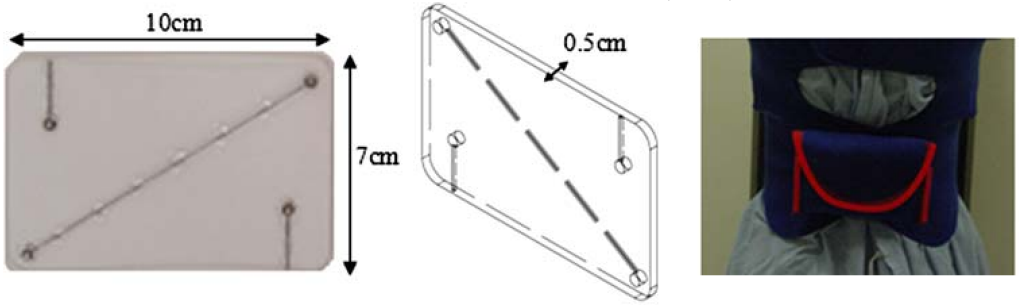


FIGURE 2.10: Small calibration object proposed by Kadoury et al. (2007b) with 4 pellets of known position.

Very recently, Kadoury et al. (2010) proposed using high level primitives, such as the 3D visual hull of the spine automatically extracted from the radiographs. This approach, while dependent of the segmentation quality, enables to avoid some of the pitfalls of the manual identification of stereo-corresponding points.

Source: adapted from Kadoury et al. (2007b)

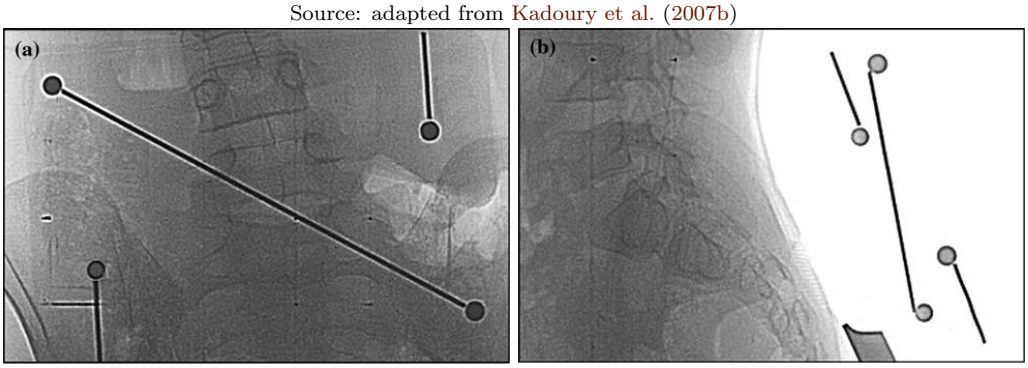


Figure 2.11: Radiographs showing the calibration object proposed by Kadoury et al. (2007b). All 4 pellets are visible in both radiographs.

Still, the same object proposed by Kadoury et al. (2007b) was used to provide an initial guess of the calibration parameters. Simulation experiments show that using high level primitives provides higher tolerance to noise, whereas *in vivo* experiments show that the method is sufficiently accurate for enabling the calculation of angular clinical indices. However, distance-based indices like the spinal length were not evaluated nor the accuracy of the 3D reconstructions. This work shows that calibration robustness may be improved using high level primitives but still relies on calibration objects for an initial guess. Unfortunately, scale recovery was not evaluated and so remains unclear if this approach provides any kind of improvement on this issue. Once more, a rotatory platform for proper patient positioning was used.

Small calibration objects have almost the same advantages as not using calibration objects at all. However, using small calibration objects results in higher costs (the cost of the calibration object), adaptations to the radiological procedure, and more importantly, the objects usually overlap anatomical structures on radiographs (as it is possible to observe in figure 2.11). In terms of accuracy, using small calibration objects is obviously advantageous, although for best results one should use large calibration apparatus.

Besides the methods described here, other solutions are used for general radiography. One standard technique is to use a small calibration sphere with known dimensions that should be aligned with the anatomical structure under examination. With only one sphere it is possible to make a rough correction of the scaling effect by using the known sphere diameter and the observed diameter on the radiograph. For instance, The (2006) was able to calibrate pelvis radiographs using one sphere that should be positioned between the patient legs and aligned

with the pelvis. To our knowledge, this technique was never applied to spine radiography since the large dimensions and irregular shape of the spine would worsen the scaling effect. Additionally, this technique only allows to quantify measures on the plane of the radiograph, and thus not enabling a 3D assessment of the spine.

2.1.4 Hybrid methods

Cheriet et al. (2007) used a novel technique to calibrate spine radiographs. It consists in mixing two kind of calibration objects: one object with absolute positioning in the radiological environment, and a vest worn by the patient with a set of calibration pellets (figure 2.12). The first object offers an absolute reference plane that helps calculating calibration parameters. On the other hand, the vest offers a set of points that follow the patient and enable to compensate undesired changes of posture between acquisitions. Moreover, the vest offers accurate stereo-corresponding points that are more reliable and easier to localise than anatomical points. In this way, the calibration algorithm only uses positioning information provided by calibration objects (the fixed object and the vest), which are more accurate and easier to detect. Thus, the authors were able to automatically localise the calibration objects using image processing techniques (Kadoury and Cheriet, 2006). Currently, only preliminary results are available indicating that this method is more robust to patient movements between radiographs comparing to large calibration apparatus, which only try to minimise patient movements using rotatory platforms. Additionally, clinical indices are more accurate than the method proposed by Kadoury et al. (2007b). To the best of our knowledge, neither accuracy measuring distances nor accuracy of geometrical reconstructions with this calibration method were published.

Source: adapted from Cheriet et al. (2007)

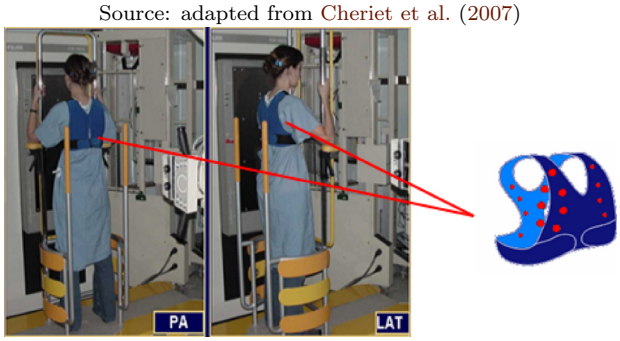


FIGURE 2.12: Calibration vest proposed by Cheriet et al. (2007) with 16 pellets.

Recently, a commercial system suggests a different approach to tackle the problem of compensating changes in posture between radiographs, which is based on the acquisition of 3D surface data simultaneously with the radiographs (Blanchard and Elbaroudi, 2008). The patient wears a vest with radiopaque objects that are also visible on the surface, more precisely at the patient backs. The light source is placed in a fixed position that enables to capture the calibration objects on both the frontal and lateral views. The 3D position of the calibration objects is acquired by structured light simultaneously with each radiograph. The 3D data is then used to determine changes in posture between the two acquisitions and, thus, to update the 3D reconstruction from the radiographs until the difference between the reconstructed calibration objects using the two acquisition techniques is negligible. No validation study was found for this system, therefore it is not possible to compare it with the previous methods. In our point of view, it is difficult to guarantee that radiographs and surface are acquired simultaneously, which may jeopardise correction.

2.1.5 Summary

In this section we reviewed spine radiography calibration. Large calibration apparatus are still the gold-standard in terms of reconstruction accuracy. Small calibration objects offer more flexibility and lower costs but are also less accurate. Nevertheless, small calibration objects offer sufficient accuracy for several clinical applications. Using calibration objects, either large or small, always have disadvantages, such as overlapping anatomical structures, costs, and changes in the protocol. Thus, calibrating spine radiographs without using calibration objects would be an interesting approach, but presently no method was able to perform accurate reconstructions. Recovering scale without calibration objects remains an open issue of research.

Novel techniques are emerging that try to compensate undesired patient movements between acquisitions by either using anatomical features of the patient spine, or calibration jackets instead of a fixed setup surrounding the patient. However, despite some of these approaches demonstrated that are sufficiently accurate for the calculation of angular clinical indices, their accuracy reconstructing the spine in 3D was not proven.

2.2 Construction of geometrical models of the spine from radiographs

After determining the calibration parameters for two (or more) radiographs of the same subject, one is able to calculate the real 3D coordinates for every point that is visible in both images (stereo-correspondent points). However, for properly describing the spine, several points are required. Knowing the location of the pedicles and endplates of each vertebra allow to analyse global, spinal and regional deformities (Stokes, 1994), although more data is required for evaluating local (vertebral) deformities. In the next subsections we will present the main approaches for constructing personalised geometrical models of the spine from radiographs. Most of the methods aim at recovering the global structure of the spine, but not all try to reconstruct vertebrae' shape.

2.2.1 Highly-supervised methods

The methods on this subsection need considerable user interaction, more precisely, several points must be manually identified for each vertebra in every radiograph in order to achieve reconstructions of the spine. Nevertheless, some of them are still used in clinical environment.

Stereo-correspondent points and kriging

The first method for constructing morpho-realistic models of the spine from radiographs was proposed by Aubin et al. (1995). This method requires 6 landmarks per vertebra per radiograph that must be manually identified by an expert. These stereo-correspondent landmarks are 6 anatomical points visible in both radiographs, i.e. centre of superior and inferior endplates, and the superior and inferior extremities of the pedicles (figure 2.13). Their 3D coordinates are found by triangulating the 2D coordinates of the manually identified points. Thus, six world coordinates are known for each vertebra, which describe the global structure of the spine. For building personalised vertebrae models the authors used a generic model for each vertebra (acquired from a single CT scan of a cadaver, figure 2.14) that was deformed using dual kriging* (Trochu, 1993) in order to fit the six landmarks.

*. Kriging is a technique for interpolating and extrapolating values for unobserved locations using information of nearby locations.

Source: adapted from Mitulescu et al. (2002)

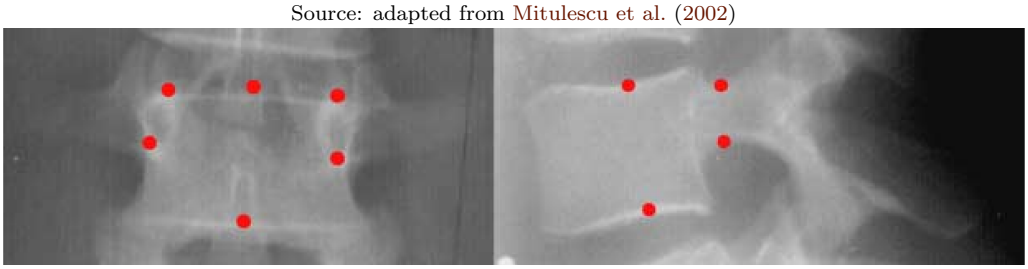


Figure 2.13: Stereo-corresponding anatomical landmarks: the points identified in one view have a correspondent point in the other view and therefore the number of points is the same in both views (pedicles are overlapped in the lateral view).

Source: adapted from Mitulescu et al. (2002)

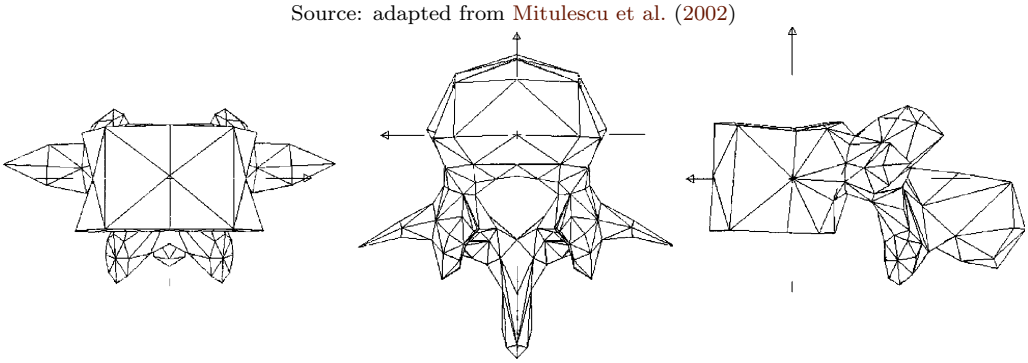


Figure 2.14: Generic model of a lumbar vertebra that is used as reference for constructing personalised lumbar vertebrae.

This method was evaluated by Aubin et al. (1997) using a cadaveric spine composed by 17 dry non-pathological vertebrae (T1–L5). The authors compared the 3D coordinates of a set of 21 landmarks per vertebra measured using a coordinate measuring machine (error: ∼0.1mm) with an equivalent set of landmarks extracted from a 3D reconstruction using the proposed method with three radiographs (PA, PA 20° and Lateral). Point-to-point errors were of 2.6 ± 2.4mm (mean ± standard deviation).

This work also studied the effect of increasing the number of landmarks on the method accuracy. For testing this, experiments where done using the full set of 21 landmarks per vertebrae per radiograph. As expected, errors decreased but because of the exaggerated time required for manually identifying all the points and because some of the points were sometimes difficult to see, this option was considered inadequate for routine utilisation (Aubin et al., 1997). The authors did not tested varying the number of landmarks for determining a good compromise between the time spent and the achieved accuracy. Moreover, when using the full set of 21 landmarks as input, the validation set is equivalent to the input set

of landmarks and, therefore, there are no landmarks left to assess the method performance in regions where no input is available. Using a larger set of validation points would enable a more realistic assessment of the method's performance recovering shape.

Adding non stereo-corresponding points to SCP and kriging

Using stereo-correspondent points with kriging is not enough for obtaining accurate 3D vertebrae reconstructions. Adding more stereo-correspondent points besides the six points proposed by Aubin et al. (1997) may lead to more accurate reconstructions, however there are several important anatomical features that are only visible in one of the radiographs (when considering the standard frontal and lateral radiographs). Having this in mind, Mitton et al. (2000) proposed enhancing the previously describe method with non stereo-corresponding points (NSCP – points that are identified in only one of the radiographs). An example of non-stereo corresponding points is shown for a lumbar vertebra in figure 2.15, although in this work the authors tested their technique in cervical vertebrae (Mitton et al., 2000).

Source: adapted from Mitulescu et al. (2002)

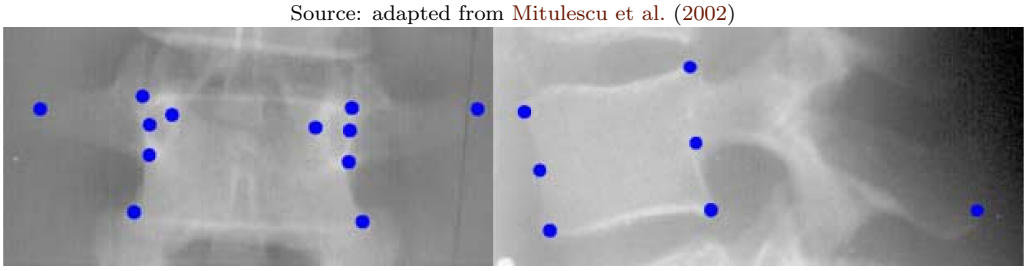


FIGURE 2.15: Non stereo-corresponding anatomical landmarks: the points marked in one view are independent of the points marked in the other view and therefore the number of points may differ from one view to the other.

The main issue with NSCPs is that it is not possible to determine the 3D coordinate by triangulation because these points appear in one radiograph only and therefore do not have a correspondent point in the other radiograph. Thus, there is no way to determine the exact 3D coordinates in the real-world for NSCPs. The only information available is that the 3D position will be somewhere in the line joining the x-ray source to the marked point in the radiograph (constraint line). To overcome this problem Mitton et al. (2000) proposed a method where a generic object (a vertebra model) is used for estimating the 3D coordinates of NSCPs. This object is composed by a set of points connected by springs with

a given stiffness that try to keep the deformed object as close as possible to the original shape (in order to maintain a generic topology). For finding a solution, an optimisation problem is defined where the goal is to minimise the total deformation, constrained to the location of the SCPs and to the constraint lines of the NSCPs. The authors used a conjugate gradient algorithm to compute the solution for the minimisation problem.

Summarising, the algorithm for constructing geometrical models of the spine has the following steps:

1. Manual identification of the SCPs (6 per vertebra per radiograph) and NSCPs (15-26 per vertebra);
2. Calibration, e.g. using DLT (see section 2.1);
3. Computation of the 3D coordinates of the SCPs using triangulation;
4. Computation of the 3D coordinates of the NSCPs using the previously described optimisation process;
5. Kriging of a generic object (with approximately 200 points) using the 3D points previously computed.

This algorithm was first applied for reconstructing dried cervical vertebrae (Mitton et al., 2000), then for dried lumbar vertebrae (Mitulescu et al., 2001) and finally for thoracic and lumbar vertebrae *in vivo* (Mitulescu et al., 2002). Results from the last study achieved reconstructions errors of 1.5 and 2.0mm (mean and RMS), which shows a considerable improvement over the previous technique. Additionally, Mitulescu et al. (2001) compared reconstructions using SCP only, SCP and NSCP, and CT scans against direct measures with a 3D scanner. The results demonstrated a considerable improvement on mean and maximum errors, and also placed reconstruction from radiographs a step closer to CT scans reconstructions (table 2.1).

Later, Bras et al. (2003) increased the detail of generic vertebrae models from $\sim$ 200 points to 2000 points (per vertebra). The original generic model had been obtained by direct measurements using a 3D scanner over a large set of dried vertebra ($\sim$ 1000 vertebrae) and was updated using a CT scan (figure 2.16). This improvement seems to reduce the reconstruction error. However, the validation technique was inferior and the quality of the radiologic equipment was superior (EOS™ system), which makes difficult comparing results. Nevertheless, the maximum reconstruction error was decreased by 2mm, notwithstanding the accuracy of the validation technique (table 2.1).

TABLE 2.1: Comparison of 3D reconstruction errors of dried lumbar vertebrae for SCP (stereo-corresponding points only), NSCP (SCP + non stereo-corresponding points), NSCP with generic models of high detail, and CT scans. It is also presented the number of landmarks needed to be manually identified per vertebra, the number of points of the geometrical model for each vertebra, the number of vertebrae tested, and the validation method.

	CT	SCP	NSCP	
Mean error (mm)	0.8	1.4	1.1	0.9
2RMS error (mm)	2.2	3.6	2.8	2.4
Maximum error (mm)	3.9	15.8	7.8	5.8
Number of landmarks	—	$6 * 2 = 12$	$6 * 2 + 19 = 31$	31
Model detail (points)	178			2000
Test set size	18			14
Validation	3D scanner (±0.2mm)			CT (±1mm)
Source	Mitulescu et al. (2001)			Bras et al. (2003)

Source: adapted from Bras et al. (2003)

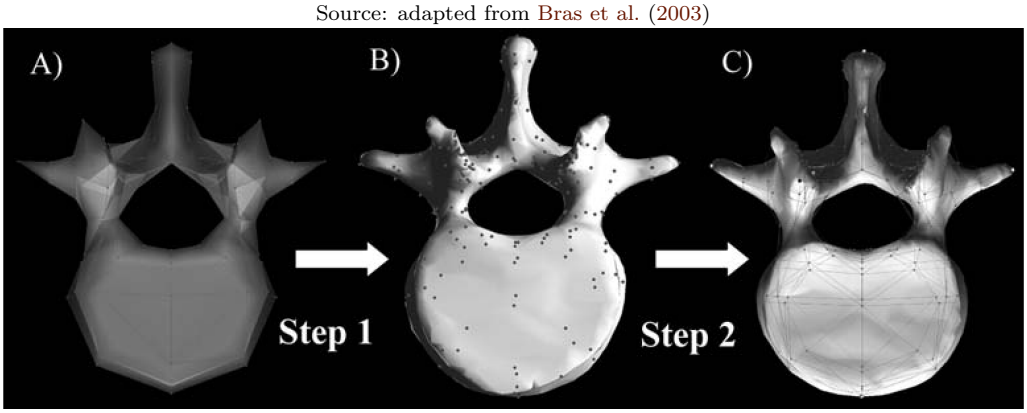


FIGURE 2.16: Refining the model of a generic lumbar vertebra; a) Previous model with 200 points; b) CT scan of a lumbar vertebra; c) New model with 2000 points.

Results presented so far for this method are only related to the accuracy recovering vertebrae shape. In another work, Dumas et al. (2004) studied the accuracy and precision determining vertebrae orientation of both NSCP and SCP methods. The authors used a set of five dried lumbar vertebrae (L1–L5) that were placed in a set of holders which had radiopaque markers for assessing the vertebrae reference orientation. This set was placed in 15 different configurations where vertebrae rotation for each axis was selected randomly. Two orthogonal radiographs were acquired for each configuration. Vertebrae rotation for the three axes was assessed first for reconstructions using 6 SCP per vertebrae and per radiograph, and then using another 14 NSCP per vertebrae, making a total of 26 points per vertebra. Once again, using NSCPs enabled achieving better results (table 2.2). The highest errors were obtained for the axial rotation with RMS

Table 2.2: Orientation errors for SCP (total of 12 points per vertebrae) and NSCP (total of 26 points per vertebrae).

	SCP errors (degrees)			NSCP errors (degrees)		
	Mean	RMS	Max	Mean	RMS	Max
Lateral	1.2	1.7	3.6	0.6	0.8	2.0
Sagittal	1.9	2.3	4.4	0.7	1.0	2.5
Axial	2.2	2.7	5.3	1.4	1.9	4.6
Source	Dumas et al. (2004)					

error equal to 1.9° and maximum error of 4.6°. Lateral and sagittal rotations were better estimated with RMS errors of approximately 1°.

NSCPs enabled to improve the quality of the geometrical reconstruction of the spine but increased the number of points that need to be manually identified. Consequently, the time technicians must spent to reconstruct a spine is also increased and varies between 2 to 4 hours, depending of the image quality and of the patient pathology (Pomero et al., 2004) . This is a factor that obviously makes difficult implementing this method in a routine clinical environment, especially because of the required human resources.

Inferring shape and pose, vertebra by vertebra

The previous methods solely rely on the manually identified points for determining vertebra shape and pose. On the other hand, the method proposed by Pomero et al. (2004) uses statistical data for decreasing the number of points per vertebra that must be manually identified. Despite the authors labeled this method as Semi-Automated (SA), we classified it as highly-supervised because it still requires an operator to identify a set of points for each vertebra being reconstructed.

The method proposed by Pomero et al. (2004) uses a data set of 3202 vertebrae, where for each vertebra is stored its geometrical model (∼200 points) and measurements of the vertebral body. Using this data set the authors were able of decreasing user-interaction to 4 points per vertebra per radiograph (making a total of 8 points per vertebra). These points correspond to the 4 corners of every vertebrae body, independently of its rotation (figure 2.17). Using these points it is possible to calculate the 3D box (more precisely the hexahedron) that encloses every vertebral body. With this box the proposed algorithm first calculates vertebrae lateral and sagittal angulation. Then, axial rotation is inferred from the

spine curvature using a statistical model. At this point, the position and orientation of every vertebra is known. For estimating the shape of a given vertebra, a set of descriptors is calculated from its enclosing box that are used as input to a linear regression model (based on the vertebrae data set) that determines a set of 21 points. Finally, once again kriging is used to deform a generic vertebra model of 2000 points using the previously determined 21 points.

Source: adapted from Pomero et al. (2004)

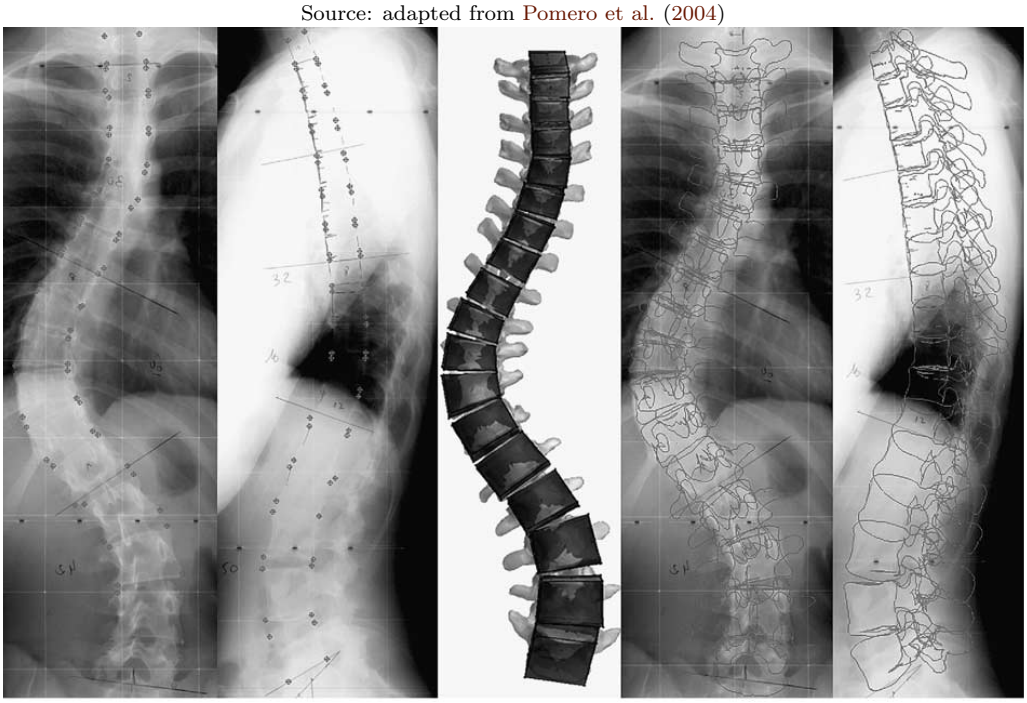


FIGURE 2.17: A less exhaustive method that needs fewer landmarks per vertebra, which are also easier to identify. From left to right: landmarks in frontal and lateral radiographs; 3D reconstruction result; projection of the 3D reconstruction to the original radiographs.

In this technique, almost every vertebra characteristic is estimated using statistical knowledge for the vertebrae level in question, ignoring adjacent vertebrae and the global shape of the spine as well as image data. For evaluating the accuracy of the proposed method, the authors used the same set of radiographs that was used by Mitulescu et al. (2002) to evaluate the NSCP technique. The accuracy of the statistical method for shape reconstruction was slightly superior and achieved a significant decrease of the maximum error, while the required user time decreased to about 20 minutes (table 2.3). However, results were not so good in terms of vertebrae orientation. Table 2.4 shows the orientation errors when compared with reconstructions using the NSCP technique for the same dataset. Unfortunately, there was no ground-truth available for vertebrae orientation be-

Table 2.3: Comparison of 3D reconstruction errors of 58 vertebrae from 14 patients for NSCP and SA (Semi-Automated method). It is also presented the number of landmarks needed to be manually identified per vertebra, and the time required to do it.

	NSCP	SA
Mean error (mm)	1.5	1.4
2RMS error (mm)	4.0	3.6
Maximum error (mm)	19.7	15.8
Number of landmarks	$6 * 2 + 19 = 31$	$4 * 2 = 8$
User time required	2–4h	$\sim$ 20m
Source	Mitulescu et al. (2002)	Pomero et al. (2004)

Table 2.4: Orientation errors for SA when using NSCP reconstructions as basis.

	SA vs NSCP Mean $\pm$ SD (degrees)
Lateral	0.2 ± 3.6
Sagital	-1.1 ± 3.7
Axial	-1.3 ± 6.4
Source	Pomero et al. (2004)

cause this was an *in vivo* study. Nevertheless, when taking into consideration the validation of the NSCP technique in terms of orientation (Dumas et al., 2004) (previously presented in table 2.1), one may easily observe that the standard deviation of the orientation error of SA itself is considerably higher than the maximum absolute error of NSCP for all three axes. Therefore, maximum orientation errors for the SA are expected to be larger than for NSCP, especially for axial rotation, which is inferred from the curvature despite the correlation between the two is not strong (Person's correlation between 0.73 and 0.77 (Stokes, 1989; Pomero et al., 2004)).

In terms of reliability, Gille et al. (2007) studied the variation of clinical parameters obtained from the geometrical reconstructions for (i) the same user and the same radiographs at different times (intra-reliability) and (ii) for different users using the same radiographs (inter-reliability). The authors concluded that the observed variations were inside the expected range when analysing mild scoliotic patients, and in fact, almost all of these measures achieved high correlations in intra- and inter-reliability analysis.

When compared to NSCP, the method proposed by Pomero et al. has comparable shape accuracy, inferior orientation accuracy (where NSCP remains the best option), and requires much less user intervention. Additionally, SA performance was very stable with radiographs of different image quality, while NSCP

performance decreased when using lower quality images. The number of landmarks was decreased and are easier to identify because they have good visibility and do not require the user to imagine the location of vertebrae regions that may be difficult to see (like in NSCP). In addition, automatically detecting the corners of vertebrae seems to be a much easier segmentation task. In fact, Deschenes and de Guise (2002) have proposed a segmentation algorithm which is used in Deschenes et al. (2003) for detecting these 4 points per vertebra per image based on a small set of points hand-marked along the spine and two line-segments also identified by the user. Unfortunately, we have not got access to this last document, and to our knowledge no results have been published about the reconstruction accuracy when combining SA from Pomero et al. (2004) with the semi-automatic point detection from Deschenes et al. (2003). Moreover, in a recent anthropometric study of exterior and interior geometry of 64 adults performed by the same institution (Bertrand et al., 2006), the more supervised version of this method was used, which possibly indicates that the more automatic method achieves inferior results or that it was not conveniently evaluated.

2.2.2 Towards unsupervised reconstruction of the spine

The methods described in this section were the first where image data was used in order to decrease user interaction and to improve vertebrae reconstruction. They are based on 2D/3D registration of statistical deformable templates on the radiographs and their goal was to minimise user-input.

Vertebrae reconstruction using statistical deformable templates

Benameur et al. proposed combining contour extraction from radiographs with *a priori* knowledge about vertebrae shape (Benameur et al., 2001, 2003). This knowledge was captured in statistical deformable templates (Cootes et al., 1992) (one template per vertebral level) that integrates a set of admissible deformations related to pathological changes observed on a representative scoliotic vertebrae population (figure 2.18). For each vertebra level, the authors used 30 normal vertebrae to define the mean vertebra shape, and 30 scoliotic vertebrae to define the variations around the mean shape.

For deforming templates, the authors used the contours of each vertebrae, which were extracted from radiographs using a popular edge-detection algorithm (Canny, 1986). Then, an optimisation problem was defined where an energy

function measuring the distance of the edges in the radiographs to the silhouette of the deformed template in the two views was minimised using a gradient descent algorithm.

Source: adapted from Benameur et al. (2003)

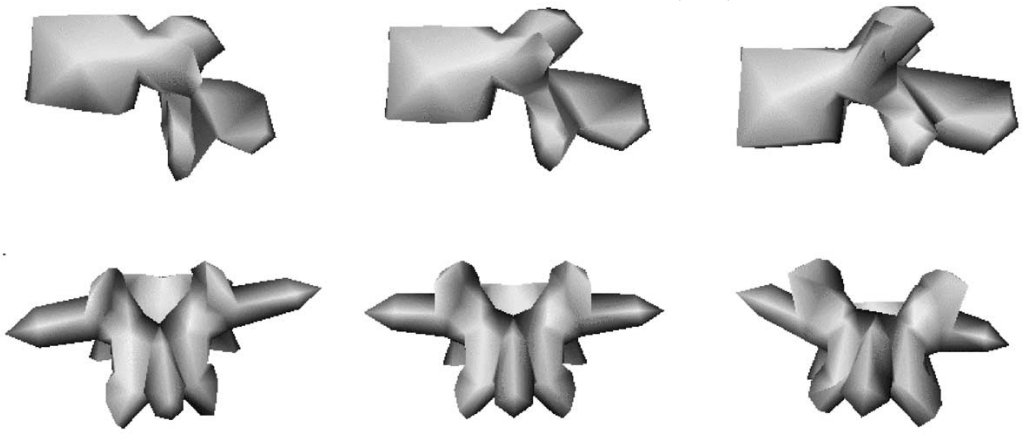


Figure 2.18: Illustration of the deformable model used in Benameur et al. (2003): mean shape for the T6 vertebrae (middle column) and variations of the first principal component (left and right columns) on the sagittal (top row) and coronal views (bottom row).

The process just described should enable to deform vertebrae templates until they fit patients' vertebrae, but in order to avoid getting trapped in local minimum it needs a good initialisation. The authors used a technique that estimates the position of 6 anatomical landmarks per vertebra (thoracic and lumbar) based on 8 points manually identified in the two radiographs (total of 16 points). This technique allowed having an initial guess of vertebrae shape and position, but during experiments it was verified that this estimation was not always good enough for escaping local minimums. The authors solved this problem by creating variations of the initial guess and thus providing several initial solutions for the optimisation algorithm. However, the success of this approach seems to be very sensitive to the parameters used for obtaining such initial solutions.

Using the described method the authors achieved results comparable with the other methods in term of vertebrae shape. These results may seem better at a first glance, but when carefully analysed it is possible to observe that for lumbar vertebrae only one or two vertebrae were validated per level (two L1, one L2, two L3, zero L4, one L5). We believe that this is a very small number of vertebrae per level, especially when each level has a specific template. Nevertheless, the authors obtained a mean error of 0.71 ± 0.06mm and maximum of 4.91 ± 0.15mm for lumbar vertebrae. Another remark is that the maximum error

is not the absolute maximum but rather an average of the maximum errors of every vertebrae reconstruction, and therefore it makes difficult comparing these results with the presented by the other authors. For thoracic vertebrae more specimen were available (reaching 10 vertebrae for the T8 level) and the mean error was 1.48 ± 0.15mm and the maximum was 7.28 ± 4.00mm. According to the authors, higher errors in the reconstruction of thoracic vertebrae are related to overlapping bone structures (e.g. ribs) that produce other contours that distract the optimisation procedure. We suppose that it was also for this reason that studies using this method were limited to vertebrae T6 to L5, since in the upper half of the thoracic vertebrae (T1 to T6) there are much more overlapping structures, especially on the lateral view.

Source: adapted from Benameur et al. (2005)

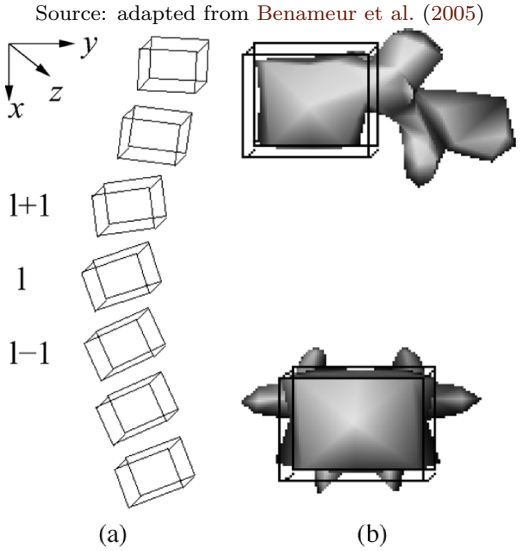


FIGURE 2.19: Deformable model of the spine used in Benameur et al. (2005) (a) and cubic template representation of a vertebra (b).

Later, Benameur et al. (2005) proposed a more unsupervised method where only 2 points per radiograph are needed (making a total of 4 points per patient). In this approach the authors used a statistical template for registering the complete spine (figure 2.19) and consequently determining an initial solution (instead of using the previous algorithm that requires 8 points per radiograph). Then, a similar process to the one used in the previous approach is applied for improving vertebrae shape and pose. This technique leads to two separate minimisation procedures solved with a exploration/selection stochastic algorithm. The accuracy was inferior when compared with the previous method using the same test set, but remains comparable with the methods from the previous section. Only

the maximum errors of thoracic vertebrae were decreased.

2.2.3 Semi-supervised spline-based methods

All the methods described in this section are based on the manual identification of the spine midline in two views (frontal and lateral) through the use of cubic splines. The spine midline is much easier and faster to identify than other anatomical features and provides a solid ground for the reconstruction procedure. Nevertheless, some of these methods require additional input data. Some of the works described here aim at capturing vertebrae location and pose as well as measuring clinical indices, rather than recovering vertebrae shape with accuracy.

Inferring 2D anatomical landmarks

To our knowledge, the first attempt of using splines for reconstructing the spine in 3D was proposed by Vaiton et al. (2004). In this work, the authors proposed creating an independent statistical model for each view (Posterio-Anterior and lateral). Each model had as input the control points of a 2D spline representing the spine midline and as output the 2D coordinates of 6 anatomical points (i.e. centre of superior and inferior endplates, and the superior and inferior extremities of the pedicles) for that view. The 3D coordinates are found by triangulating the two independent inferences from the two views.

Inferring was based on updating a mean model based on properties of the spline (e.g. its length). For couping with the different types of deformities several models had to be created according to the type of curvature of the spine and its severity, besides the view of the radiograph (figure 2.20). Additionally, the method relied on a fixed number of control points, therefore a study was performed to determine the best number of control points for each one of the models. A total of 791 radiographs were used for creating the statistical models. *In vivo* experiments shown errors of 3.0 ± 2.3mm for the endplates and 4.4 ± 2.5mm for the pedicles (mean $\pm$ SD of RMS reconstruction errors) when compared with the reference method (SCP – section 2.2.1). Despite these errors, differences between the two methods measuring clinical indices were not significantly different (Wilcoxon test, $p \leq 0.05$).

This method presents several flaws. For starting, inferences are done at 2D when the goal is to recover 3D data. The estimation of the two views are in-

Source: adapted from Vaiton et al. (2004)

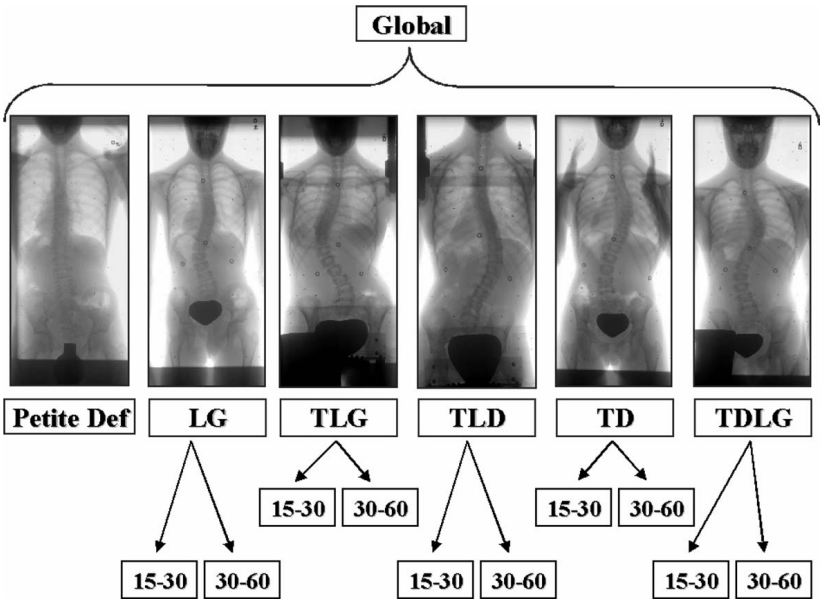


FIGURE 2.20: Classification of curvature type and severity taken into account by the method proposed by Vaiton et al. (2004). (For each curvature and each view the method uses a different model with a pre-determined number of control points for the input splines.)

dependent and then triangulated, which may jeopardise the reconstruction. Additionally, 2D data depends of the positioning of both patient and x-ray source during the examination. Therefore, this method is only suitable for radiographic protocols where the setup is fixed, which was the case. Finally, the method forces the operator to identify a specific number of control points according to the type and severity of curvature, as well as the view, which is not practical.

Interpolating 3D vertebrae size, pose and location

Dumas et al. (2008) proposed using the spine midline for restricting computation as well as user interaction. After manually identified the spine midline, users have to manually adjust the scale and location of the first and last vertebrae (C7 and L5), constrained to the splines, until the projected silhouettes on both views best match the radiographs. After this, the size and location of the remaining vertebrae are linearly interpolated along the spline. Axial rotation is calculated from the lateral shift of the spine on the frontal radiograph. Then, an optimisation of vertebrae location takes place that tries to adjust inter-vertebra space according to the vertebral level. Finally, the operator may manually refine

the reconstruction, this time with no constrains.

The study presented by Dumas et al. (2008) does not include validation of vertebrae shape since the goal of the method was to capture the global shape of the spine. The accuracy of the method was evaluated by comparison with the SCP method (section 2.2.1) in terms of location and orientation of vertebrae. Additionally inter-operator precision was also measured. Accuracy results for mild scoliotic patients show RMS_{SD} errors from 0.9mm on the anterior-posterior axis, to 2.3mm on the proxima-distal axis, while orientation ranges from 1.3° on the lateral rotation to 3.2°. These results show that the method has difficulties finding the vertebrae location along the spine midline, as well as rotation about the spine midline (axial rotation). Despite the method tries to capture vertebrae size, no evaluation directly evaluated this feature. Additional, clinical indices were not evaluated.

The amount of interaction is still considerably high, although average reconstruction time was decreased to about 5 minutes on patients with mild deformities of the spine. The method does not use statistical data to infer vertebrae properties. Rather than that, it relies on linear interpolation and on *ad hoc* optimisation to infer the reconstruction of the spine. The operator is responsible for ensuring valid reconstructions by manual adjustments, making the method very user-dependent.

Inferring vertebrae 3D shape, pose and location

In Humbert et al. (2007, 2009), the splines, as well as additional user input (i.e. location, size and orientation of 6 predefined endplates, and position and shape of the apical vertebra and two end vertebrae), are used as predictors for inferring vertebrae position and pose along the splines, as well as vertebrae shape. Based on the input data, the method starts by estimating eight descriptors for each vertebra, namely the depth, width and position of each vertebral endplate along the spinal curve. This is done using multi-linear regression based on 175 spine reconstructions obtained with the method proposed in Pomero et al. (2004). Then, the 3D coordinates of 19 points per vertebrae are estimated using multi-linear regression based on a database of 1628 digitalised vertebrae. Vertebrae shape is recovered using these points to deform a highly detailed vertebral model of ∼2000 points. (Bras et al., 2003). A trained operator requires about 2.5 minutes for performing a fast reconstruction, although it is possible to refine reconstructions, increasing the interaction time to 10 minutes.

In this study, the authors only validated the accuracy of vertebrae shape by comparison with CT for a total of 40 vertebrae (from T11 to L5). For the fast reconstruction the mean absolute point-to-surface error was of 1.3mm and for the refined the error decreased to 1.0mm. Again, validation of vertebrae shape was limited to the lower spine, and no conclusions can be made concerning the method ability to reconstruct vertebrae above T11. In addition, the study evaluated inter and intra-observer variability for vertebrae location, orientation as well as clinical indices. This study shown that the method has superior precision to the method proposed by Dumas et al. (2008) locating vertebrae as well as finding their orientation. However, no accuracy study was done to quantify how distant the metrics provided by this method are from the reality.

Despite reducing interaction time to about 2.5 minutes, this method still requires considerable user-interaction. Additionally, linear regression was used to infer several features that have a non-linear relation with the input. Finally, like discussed on section 2.2.1, the method that was used for constructing the spine database (Pomero et al., 2004) is not able of properly capturing axial rotation, which suggests that the method proposed by Humbert et al. suffers from the same problem. An accuracy study would help clarifying this issue. Nevertheless, the accuracy of vertebrae shape for the refined reconstructions seems to be comparable or superior to image-based methods.

Statistical and image-based reconstruction

Very recently, Kadoury et al. (2009a) proposed using a statistical approach for finding an initial reconstruction of the spine, which was then refined, vertebra by vertebra, using both statistical and image data. Local Linear Embedding (LLE) Roweis and Saul (2000) was used for mapping 3D splines with normalised height to a lower-dimensional space, which was then used to infer the spine reconstruction using Support Vector Regression (SVR). A total of 732 spine reconstructions were employed for computing both the LLE and one SVR per output feature. Since this step predicted the 3D position of 6 anatomical points for each one of the 17 vertebrae (T1 to L5), a total of 306 SVRs had to be trained. The reconstruction was then refined, vertebra by vertebra, with a bundle adjustment that optimised the edge image alignment, and the epipolar geometry, constrained by the likelihood of the 2D morphology of the current solution on each radiograph. The edge image alignment is calculated using (a) the 2D silhouette of a 3D model created from a single CT scan and (b) the gradient of the image. Since all ad-

justments are made in 2D, an epipolar geometry error was included to enforce valid correspondences between particular landmarks of the two images.

Average computation time by itself was of 2.4 minutes per reconstruction on an 1.9 GHz Intel PC, and user-interaction time for identifying the splines was not reported. *In vivo* reconstruction errors for patients with mild scoliosis were compared to the SCP method (section 2.2.1) and errors were of 2.2 ± 0.9mm for the endplates, and 2.0 ± 1.5mm for the pedicles (mean $\pm$ SD of RMS errors of each reconstruction). A shape validation was also performed by comparison with MRI for a total of 9 vertebrae on the lower half of the spine. Mean point to surface errors after rigid registration were of 1.2 ± 1.1mm for lumbar vertebrae (L2 to L3) and 1.1 ± 0.8mm for thoracic (T10 to T12). Unfortunately, the number of vertebrae used for this validation as well as its range in terms of vertebral level is very small to make conclusions about the capabilities of the proposed method for recovering vertebral shape. Nevertheless, concerning shape recovery of the lower vertebrae, this method does not seem to outperform the method proposed in Benameur et al. (2003), which is also image-based, and neither Humbert et al. (2009) that is based on statistical inference and manual adjustments. This study did not evaluate any clinical indices.

In the statistical approach of this study, the spine midline (with normalised scale) is the only predictor of the shape of the spine. While this is acceptable, there may be a range of spine shapes for the same spline curve. Additionally, LLE may produce inaccurate predictors for splines that are not sufficiently well sampled (van der Maaten et al., 2009). This obliges the use of very large databases (in this case, with 732 spines) that may be impractical to collect, but even then, there might be deformations that do not have enough closely related neighbours to generate accurate results. Concerning the regression phase, using a set of independent SVRs does not ensure plausible reconstructions of the spine since each output feature is trained independently and, thus, the longitudinal relation between vertebrae is not taken into account. Regarding the refining phase, vertebrae are adjusted one by one without having into account inter-dependencies between adjacent vertebrae. Moreover, most of the statistical knowledge included in this phase is two-dimensional, making it dependent of the clinical setup, and may not conveniently describe the 3D morphology of vertebrae. Nevertheless, despite these issues, and to the best of our knowledge, this remains the semi-supervised method that shows more close results to highly-supervised methods concerning the shape of the complete spine.

2.2.4 Summary

Highly-supervised techniques, such as SCP and NSCP, are the most studied and validated among all methods for reconstructing the spine. However, they require an expert for identifying an extend set of anatomical points, which makes of them very time-consuming (2 – 4 hours per patient) and error-prone. Naturally, several approaches tried to decrease user-interaction. In (Pomero et al., 2004) an alternative that remained highly-supervised enabled to decrease reconstruction time to about 20 minutes; however, axial rotation, which is statistically inferred, seems not to be properly captured. Recently semi-supervised methods based on the identification of the spine-midline trough cubic splines enabled to reduce user-interaction even-further. Interaction time is decreased to 2.5 min with the method proposed by Humbert et al. (2009). However, neither accuracy of vertebrae location nor of clinical indices were evaluated, but vertebrae shape for the lower spine seems to be properly captured. Image-based methods, such as the ones proposed by Benameur et al. (2005) and Kadoury et al. (2009a), also only show results concerning vertebrae shape for the lower spine where the images are clearer. Nevertheless, the method proposed by Kadoury et al. (2009a) show promising results for patients with mild scoliosis when compared with a fully-supervised method.

Currently, to the best of our knowledge, no method is capable of performing accurate reconstructions of the spine automatically, or with very limited user input. Additionally, no semi-supervised method has shown both accurate reconstructions while demonstrating clinically applicability by determining clinical indices with no significant differences from fully-manual approaches. Finally, despite semi-supervised methods enabled to reduce reconstruction time to minutes (instead of hours), the level of user interaction or computation time are still high for clinical routine use.

Table 2.5: Synthesis of studies evaluating the accuracy to recover the global shape of the spine (without vertebra by vertebra alignment with the reference data). (pt/v – points per vertebra; pt/P – points per patient; spl/P – splines per patient; spl^+/P – splines per patient, plus additional input.)

Reference	Calibration	Sample	Validation	User Input	Error (mm)
		Calibration studies (section 2.1)			
Aubin et al. (1997)	Dansereau and Stokes (1988)	*In vitro*, 1 spine	3D scanner ($\pm$0.1 mm)	$12pt/v$	Plates: Mean=1.5 SD=0.7 Pedicles: Mean=1.2 SD=0.7
Cheriet et al. (1999b) *	Cheriet et al. (1999b)	*In vivo*, unknown number of patients	Unknown	$12pt/v$	RMS: 2.5mm
Cheriet et al. (2007) ‡	Cheriet et al. (2007)	*In vivo*, 20 scoliotic patients	Dansereau and Stokes (1988)	$12pt/v$	Cobb angle: $0.3 \pm 0.428°$ Frontal balance: $0.15 \pm 0.151°$ Sagittal balance: $0.37 \pm 0.259°$
Kadoury et al. (2007a) ‡	Kadoury et al. (2007a)	*In vivo*, 60 patients (51 with scoliosis)	Cheriet et al. (2007)	$12pt/v$	Cobb angle: $0.29 \pm 0.15°$ Frontal balance: $0.53 \pm 0.51°$ Sagittal balance: $1.45 \pm 1.03°$ Spine Length: 14.19 ± 8.00mm
Kadoury et al. (2007b) ‡	Kadoury et al. (2007b)	*In vivo*, same as (Kadoury et al., 2007a)	Cheriet et al. (2007)	$12pt/v$	Cobb angle: $0.24 \pm 0.33°$ Frontal balance: $0.75 \pm 0.51°$ Sagittal balance: $2.62 \pm 1.65°$ Spinal length: 2.05 ± 1.03mm
Kadoury et al. (2010) ‡	Kadoury et al. (2010)	*In vivo*, same as (Kadoury et al., 2007a)	Cheriet et al. (2007)	*none*	Cobb angle: $0.31 \pm 0.26°$ Frontal balance: $0.97 \pm 0.53°$ Sagittal balance: $1.88 \pm 1.15°$
		Reconstruction studies (section 2.2)			
Vaiton et al. (2004)	Dansereau and Stokes (1988)	*In vitro*, 49 patients	Aubin et al. (1995)	$2spl/P$	Plates: Mean=3.0 SD=2.3 Pedicles: Mean=4.4 SD=2.5
Dumas et al. (2008)	Blanchard and Elbaroudi (2008)	*In vivo*, 11 scoliotic patients	Aubin et al. (1995)	$2spl^+/P$	Vert. Location RMS_{SD} = (0.9, 1.1, 2.3) Vert. Rot. RMS_{SD} = (1.3°,2.0°,3.2°)
Kadoury et al. (2009a)	Cheriet et al. (2007)	*In vivo*, 20 scoliotic patients	Aubin et al. (1995)	$2spl/P$	Plates: Mean=2.2 SD=0.9 Pedicles: Mean=2.0 SD=1.5

*Original article not available; results from Novosad et al. (2004). ‡Geometrical model accuracy not evaluated, only clinical indices; some of these are presented.

TABLE 2.6: Synthesis of studies evaluating the accuracy of vertebrae shape reconstruction. Generally, each reconstructed vertebra is aligned locally before computing point-to-surface errors. (pt/v – points per vertebra; pt/P – points per patient; spl/P – splines per patient; spl^{+}/P – splines per patient, plus additional input.)

Reference	Calibration	Sample	Validation	User Input	Error (mm)
Aubin et al. (1997)	Dansereau and Stokes (1988)	*In vitro*, 17 vertebrae (T1-L5, from 1 dried spine)	3D scanner (±0.1 mm)	$18pt/v$	Mean=2.6 SD=2.4
Delorme et al. (2003)	Dansereau and Stokes (1988)	*In vivo*, 60 vertebrae (T6-L3, from 14 scoliotic patients)	CT (±0.1 mm)	$12pt/v$	Mean=3.3 SD=3.8
Mitton et al. (2000)	Dansereau and Stokes (1988)	*In vitro*, 18 vertebrae (upper cervical): 6 C0 + 6 C1 + 6 C2	3D scanner (±0.2 mm)	$33-40pt/v$	C0: Mean=1.9 RMS=2.7 Max=16.0 C1: Mean=1.0 RMS=1.2 Max=4.3 C2: Mean=0.8 RMS=1.0 Max=3.9
Mitulescu et al. (2001)	Dansereau and Stokes (1988)	*In vitro*, 18 vertebrae (lumbar)	3D scanner (±0.2 mm)	$31pt/v$	Mean=1.1 2RMS=2.8 Max=7.8
Mitulescu et al. (2002)	Dansereau and Stokes (1988)	*In vivo*, 58 assorted vertebrae (from 14 scoliotic patients)	CT (±1 mm)	$31pt/v$	Mean=1.5 2RMS=4.0 Max=19.7
Dumas et al. (2003)	Dumas et al. (2003)	*In vitro*, same as Mitulescu et al. (2001)	3D scanner (±0.2 mm)	$31pt/v$	Mean=1.2 RMS=1.6 Max=6.4
Bras et al. (2003)	Dubousset et al. (2005)	*In vitro*, 36 vertebrae (inferior cervical, thoracic and lumbar)	CT (±1 mm)	$31pt/v$	C: Mean=0.7 2RMS=1.7 Max=4.2 T: Mean=0.9 2RMS=2.4 Max=5.6 L: Mean=0.9 2RMS=2.4 Max=5.8
Pomero et al. (2004)	Dumas et al. (2003)	*In vivo*, same as (Mitulescu et al., 2002)	CT (±1 mm)	$8pt/v$	Mean=1.4 2RMS=3.6 Max=15.8 [12.7] [§]
Benameur et al. (2003)	Dansereau and Stokes (1988)	*In vivo*, 51 thoracic (T6-T12) and 6 lumbar vertebrae (from 13 patients)	CT (±1 mm)	$16pt/P$	T: Mean=0.71 ± 0.06 Max=7.28 ± 4.00 L: Mean=1.48 ± 0.27 Max=4.91 ± 0.15
Benameur et al. (2005)	Dansereau and Stokes (1988)	*In vivo*, same as Benameur et al. (2003)	CT (±1 mm)	$4pt/P$	T: Mean=1.30 ± 1.32 Max=5.00 ± 1.06 L: Mean=1.46 ± 1.47 Max=5.95 ± 1.20
Humbert et al. (2009)	Dubousset et al. (2005)	*In vivo*, 40 lumbar and thoracic vertebrae (T11-L5, from 11 patients)	CT (±1 mm)	$2spl^{+}$/P	Fast: Mean=1.3 2SD=3.6 Refined: Mean=1.0 2SD=2.7
Kadoury et al. (2009a)	Cheriet et al. (2007)	*In vivo*, 3 lumbar and 5 thoracic vertebrae (T10-L3, from 2 patients)	MRI (±1 mm)	$2spl/P$	T: Mean=1.1 SD=0.8 L: Mean=1.2 SD=1.1

[§]Detected one incomplete vertebra in the validation scan; real maximum error inside brackets.

CHAPTER 3

Geometrical calibration of biplanar radiography

This chapter presents a novel calibration method for general biplanar radiography that extends a well known technique by including a distance on the geometrical model that can be easily measured on site (Moura et al., 2008a,b). This extension, coupled with an appropriated optimisation algorithm, enables achieving accurate 3D reconstructions while minimising the need for calibration objects. The chapter starts with a brief explanation of the goal of geometrical calibration followed by an analytical description of the geometry involved on the formation of biplanar radiographs. Then, the proposed calibration method is described, including detail about the algorithms that were used for solving the problem. Finally, the method evaluation is presented and results are discussed.

3.1 Goal of the geometrical calibration

On biplanar x-ray systems, two radiographic images are acquired that represent two different views of the same 3D target. Figure 3.1 illustrates a biplanar radiographic exam where the x-ray system stays stationary while a subject is repositioned from one acquisition to the other for providing different views of the

trunk. Several geometrical and physical parameters influence the image formation process affecting the shape of the target on the radiograph as well as the intensity values that are captured. When geometrically calibrating radiographs we are interested on the parameters that affect the shape, such as the position and orientation of the patient in relation to the system, and the distance between the x-ray source and the x-ray detector. These parameters change from exam to exam according to the anatomical structure being analysed (e.g. hand vs spine) and to the patient anatomy (e.g. child vs adult).

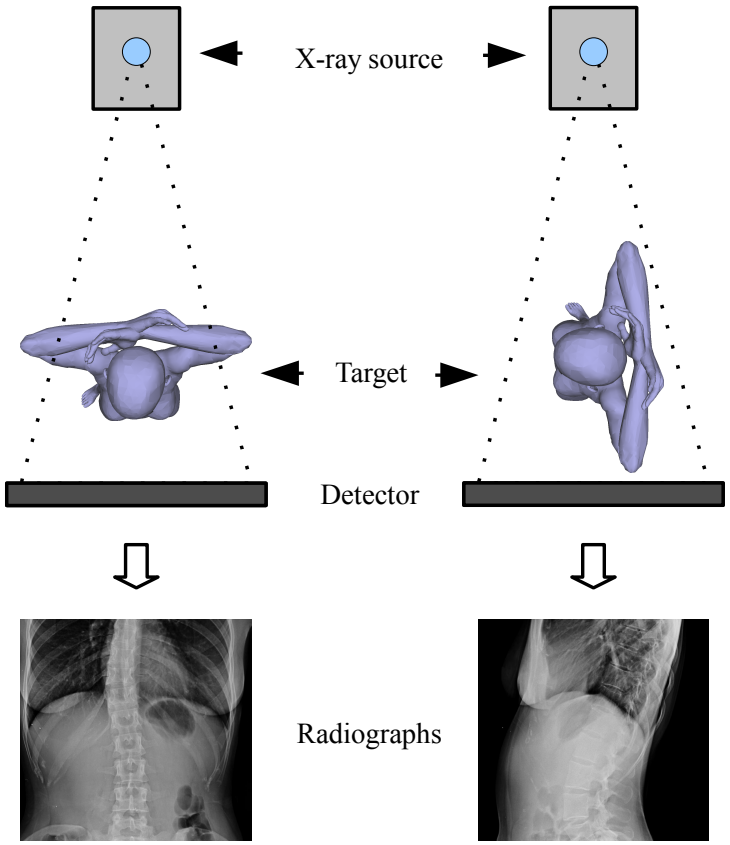


Figure 3.1: Illustration of a biplanar radiographic exam on a convectional system showing the subject's positioning on the two acquisitions while the system remains stationary (top view).

The goal of the calibration process for a single planar radiograph is to find the geometrical parameters that allows to compute the projection of 3D points in world coordinates to 2D points on the radiograph. In biplanar radiography, two radiographs are available showing two different views of the same target. Knowing the geometrical parameters of the two radiographs makes also possible to calculate

3D points in world coordinates from their 2D locations on the radiographs using a technique known as triangulation (Hartley and Sturm, 1997). Thus, calibrating biplanar radiographs allows to calculate 3D data from 2D data visible on the two views.

3.2 The geometry of biplanar radiographs formation

Image formation, including the formation of planar radiographic images, is generally modelled by a transformation that maps 3D points in the real world to 2D points in the image. Analytically, the projection of a 3D point (in homogenous coordinates) to two planar images is described by the following equation:

$$\begin{bmatrix} w_i \cdot u_i \\ w_i \cdot v_i \\ w_i \end{bmatrix} = M_i \cdot \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \quad \text{for } i = 1, 2 \tag{3.1}$$

where for each view i, M_i is a 3×4 matrix that describes the projection of the 3D point (X, Y, Z) into image coordinates (u_i, v_i) subjected to a scaling factor w_i. Matrix M_i is called calibration matrix. The goal of the calibration procedure is to determine the calibration matrix for each view.

For flat x-ray detectors, M may be modelled using the perspective projection model (Lepetit and Fua, 2005):

$$M_i = \underbrace{\begin{bmatrix} f_i/s_u & 0 & u_{p_i} & 0 \\ 0 & f_i/s_v & v_{p_i} & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{K_i} \cdot \underbrace{\begin{bmatrix} \mathcal{R}_i & t_i \\ 0_3^T & 1 \end{bmatrix}}_{G_i} \quad \text{for } i = 1, 2 \tag{3.2}$$

where K_i is a 3×4 matrix representing the intrinsic parameters of the system, and G_i is a 4×4 matrix that represents the geometrical transform that aligns the target's coordinate system with the coordinate system of the x-ray source. This model assumes that there is no pixel skew (the axes of the image are perpendicular) and no distortion that is typically present in cameras due to the use of lens (Forsyth and Ponce, 2003). These simplifications are valid since we are modelling planar detectors. Matrix K_i is calculated using:

- f_i – the focal distance, i.e. the distance in real world units (e.g. mm) between the x-ray source and the detector on radiograph i (figure 3.2);

- (s_u, s_v) – the sampling pitch of the detector for each axis of the image, which is known and constant for a given system and represents the resolution of the radiographs (e.g. in mm/pixel);
- (u_{p_i}, v_{p_i}) – the principal point, i.e. the 2D projection of the x-ray source in the radiographic image i, measured in pixels (figure 3.2).

The geometrical transform G_i is defined by a 3×3 rotation matrix $\mathcal{R}_i$ that represents the target's orientation, and a translation vector $(t_{x_i}, t_{y_i}, t_{z_i})^T$ that represents the target's position in relation to the x-ray source, with 0_3^T representing $(0, 0, 0)$.

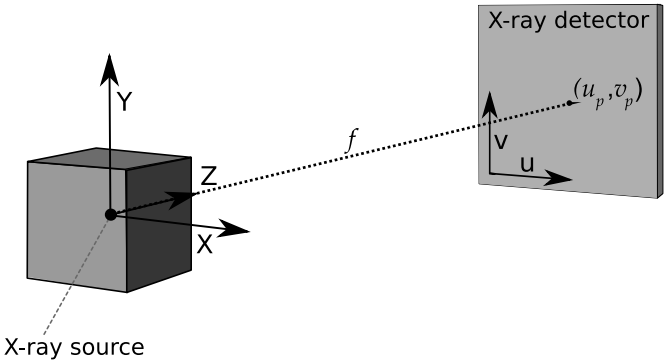


Figure 3.2: Coordinate system of reference and representation of the intrinsic parameters.

Thus, the goal of finding the calibration matrix M_i may be replaced by finding the intrinsic parameters as well as $\mathcal{R}_i$ and t_i. Still, $\mathcal{R}_i$ is a 3×3 matrix with only 3 degrees of freedom (DOF) and, therefore, it may be represented in a more compact way for better numerical stability as well as to avoid placing validity constrains when estimating pose. Typical representations of the rotation matrix are Euler angles, quaternions, and axis/angle (Schmidt and Niemann, 2001; Lepetit and Fua, 2005). Euler angles are a common representation of rotation where three scalars define three elemental rotations that are applied sequentially. This representation was used in previous work of calibration of biplanar radiography of the spine (e.g. Cheriet et al. (1999a); Kadoury et al. (2007a,b)). However, it was shown by Hornegger and Tomasi (1999) that they introduce numerical sensibility to the problem of pose estimation and therefore should be avoided. Additionally, this parametrisation may originate gimbal lock singularities where one DOF is lost, i.e. two of the three Euler angles belong to the same DOF (Lepetit and Fua, 2005; Vince, 2006). On the other hand, both quaternions, and axis/angle are numerically stable with equivalent performance in problems of rotation estimation (Schmidt and Niemann, 2001). We choose to use the axis/angle representation

because it has a more compact representation with a number of parameters equal to the number of DOF of the rotation matrix and thus it does not require validity constrains when used in optimisation processes.

The axis/angle form represents an arbitrary rotation as a rotation around an axis $a \in \mathbb{R}^3$ by an angle $\theta \in \mathbb{R}$ (in radians). Both these data are parametrised in a vector $\omega \in \mathbb{R}^3$ with:

$$\omega = \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \end{bmatrix}, \quad a = \frac{\omega}{\|\omega\|}, \quad \theta = \|\omega\|. \tag{3.3}$$

Vector ω can be converted to a rotation matrix using Rodrigues' formula (Lepetit and Fua, 2005):

$$\mathcal{R} = e^{\widehat{\omega}} = I_3 + \frac{\sin\theta}{\theta}\widehat{\omega} + \frac{1-\cos\theta}{\theta^2}\widehat{\omega}^2 \tag{3.4}$$

with

$$\widehat{\omega} = \begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix}, \quad I_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{3.5}$$

Representing rotations in axis/angle form allows to redefine the goal of the calibration procedure for biplanar radiographic systems as finding the calibration parameters ξ_i:

$$\xi_i = (f_i, u_{p_i}, v_{p_i}, t_{x_i}, t_{y_i}, t_{z_i}, \omega_{x_i}, \omega_{y_i}, \omega_{z_i}) \quad \text{for } i = 1, 2. \tag{3.6}$$

This compact representation uses nine parameters per view that can be freely manipulated to generate calibration matrices without requiring constrains for enforcing validity.

3.3 General calibration method

Despite the problem of finding the calibration parameters in biplanar radiography is similar to camera calibration, there are several constrains in biplanar radiography that limit the range of choice of algorithms. Several approaches from the Computer Vision community are based on keeping the intrinsic parameters of the cameras constant, which allows to find them offline and, thus, reducing the problem to pose estimation (Lepetit and Fua, 2005). However, in conventional radiography, the distance between the x-ray source and the x-ray detector is

not constant, and the location of the principal point may also vary from exam to exam and, thus, it is not possible to determine these parameters offline unless the system has a fixed configuration. Several approaches have tackled the problem of determining the intrinsic parameters online but require three or more views (e.g. Faugeras et al. (1992); Hartley (1994); Triggs (1997); Bougnoux (1998); Pollefeys et al. (1999)), while we only have two. Additionally, most of the approaches are only concerned with the relative pose of up to scale reconstructions (Hartley and Zisserman, 2003) while in clinical applications it is important to recover the absolute orientation for enabling to calculate several clinical indices as well as the scale for accurately measuring Euclidian distances. Finally, other approaches make use of patterns/objects in the images with known 3D geometry (e.g. Abdel-Aziz and Karara (1971); Zhang (2000)), which we want to avoid.

A well known approach for capturing the scene geometry without requiring calibration objects of know geometry is minimising the squared difference between a set of point-matches identified on the two views and their analytical projections (Hartley and Zisserman, 2003). We will call this difference the retro-projection error, although other terms may also be found in the literature, such as re-projection error or back-projection error. Such approach was previously used for the calibration of biplanar radiographs, e.g. Cheriet and Meunier (1999); Cheriet et al. (1999a, 2007); Kadoury et al. (2007a,b). The minimisation of the retro-projection error can be formulated as a least-squares minimisation:

$$\min_{\xi_1^*,\xi_2^*} \left(\sum_{i=1}^{2} \sum_{j=1}^{n} \left\| p_{ij} - prj\left(\xi_i, tri\left(\xi_1, \xi_2, p_{1j}, p_{2j}\right)\right)\right\|^2 \right), \tag{3.7}$$

where n is the number of point-matches, p_{ij} is the j^{th} 2D point identified on radiograph i, prj is the 2D projection of a 3D point as defined in equations 3.1 and 3.2, tri is a triangulation operation that calculates the 3D coordinates for a given point-match (Hartley and Sturm, 1997), and ξ_i^* are the optimised parameters for the i^{th} radiograph.

Calibration is accomplished using a standard nonlinear least-squares minimisation algorithm. This class of algorithms needs an initial solution for the calibration parameters, which are then iteratively updated towards minimising the sum of squared distances between the identified and retro-projected points. These points are usually anatomical landmarks that are manually identified or calibration objects that may be automatically located.

One of the main problems of this approach is guaranteeing a proper initiali-

sation of the geometrical parameters to avoid convergence to a local minimum. In Cheriet et al. (1999a) the authors experimented performing the calibration without calibration objects. The geometrical parameters were initialised with the mean values of previous calibrations using a large apparatus, and the retro-projection error was calculated using anatomical landmarks of the spine. However, calibrations were inaccurate and the authors concluded that calibration objects are required for proper 3D reconstructions. In Kadoury et al. (2007b) a small calibration object was proposed for obtaining an approximation of the calibration parameters and for correcting scale. Here, a different approach is proposed that is based on using a distance measuring device for determining some of the geometrical parameters, namely the focal distance, and the distance between the x-ray source and the target. The hypothesis is that using such device would allow either to eliminate the need for calibration objects or to simplify their geometry for minimal impact on the content of radiographs.

3.4 Estimating the calibration parameters using a distance measuring device

The motivation for using a distance measuring device is that in radiography, in general, the distance of the x-ray source to the target structure under examination varies from examination to examination, depending of the size of the region of interest. Therefore, focal distance (f), and the distance between the target structure and the x-ray source (t_z) are unknown. We propose attaching a distance measuring device, e.g. a rangefinder, to the x-ray source for helping to calculate these parameters (figure 3.3).

The device is only capable of measuring the distance between the x-ray machine and the table (d_m). In order to calculate f, two more parameters need to be determined:

- d_s – the distance from the x-ray source to the plane of the x-ray device where the x-rays come out;
- d_d – the distance from the table to the x-ray detector.

These parameters are fixed for a given imaging system but may be difficult to measure directly with accuracy. Therefore, we propose finding them indirectly using a pre-calibration procedure that only needs to be executed once for a given system.

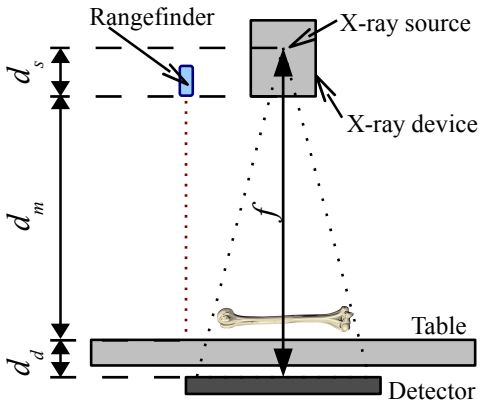


Figure 3.3: Illustration of a conventional radiographic imaging system with a rangefinder attached (top view). The x-ray table is in upright position. Distances d_d and d_m are constant for a given system, while distance d_m varies from exam to exam.

3.4.1 Pre-calibration procedure

The pre-calibration procedure here proposed aims at determining parameters d_s and d_d for a given radiographic system (figure 3.3). This procedure includes acquiring several radiographs of a calibration phantom, while varying and recording the distance d_m with, for instance, a laser rangefinder. Then, for each radiograph, the relation between d_s and d_d is determined. Finally, crossing results from all the radiographs enables to determine the values of d_s and d_d for the system.

The phantom is a radiopaque planar grid with known geometry (figure 3.4). Using a phantom with these properties allows to apply imaging processing techniques especially developed for detecting chequerboard patterns, e.g. Soh et al. (1997); Bradski (2000); Wang et al. (2007). Such techniques allow to rapidly and accurately detect the 2D position of the corners that make the chequerboard.

Let the coordinate system of reference be located at the x-ray source (figure 3.2). The pre-calibration procedure starts by placing the phantom on the table of the imaging system at a known position and orientation, i.e. centred ($t_x \simeq 0$ mm, $t_y \simeq 0$ mm), and parallel to the detector with no tilt ($\omega_x \simeq 0$, $\omega_y \simeq 0$, $\omega_z \simeq 0$, in radians). This positioning is easy to accomplish and enables to have a good initial guess for most of the calibration parameters. The x-ray source should aim to the centre of the x-ray detector and, thus, parameters (u_p, v_p) are initially set to half the size of the radiograph (in pixels). Finally, parameters f and t_z may be redefined in function of d_s, d_m and d_d:

$$f = d_s + d_m + d_d, \tag{3.8}$$

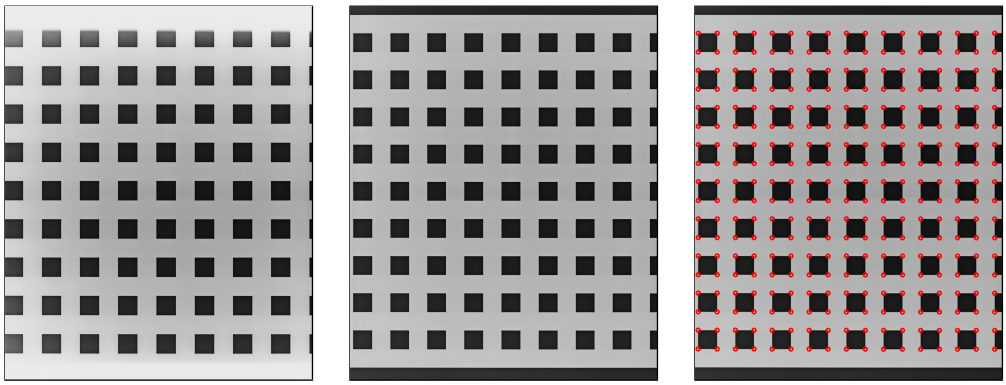


FIGURE 3.4: Radiographs of the phantom for the pre-calibration procedure acquired with $d_m = 506$mm (left) and $d_m = 909$mm (middle), and illustration of the corners extracted by a semi-automatic method (right).

$$t_z = d_s + d_m - o_t/2, \tag{3.9}$$

where o_t is the thickness of the phantom.

Since d_m is known, finding the calibration parameters that minimise the projection errors of the 3D geometry of the phantom would allow to recover d_s and d_d. However, since the phantom is planar and because of its positioning, there is an infinite range of solutions for d_s and d_d that enables to accomplish optimal 2D projections of the 3D geometry of the phantom. This happens because, for a given radiograph, increasing one of these parameters may be compensated by increasing the other. However, it is possible to capture the relation between the two parameters. This may be accomplished by fixing one of the parameters and finding the other numerically. In particular, fixing d_s allows to determine the optimal value of d_d that minimises the projection errors of the phantom geometry. Since d_m and d_s are given, the set of parameters to optimise is:

$$\zeta = (d_d, u_p, v_p, t_x, t_y, \omega_x, \omega_y, \omega_z). \tag{3.10}$$

Finding the optimal values for ζ may be formulated as a least-square minimisation problem:

$$\min_{\zeta^*} \left(\sum_{i=1}^{n} \|p_i - prj(\zeta, d_m, d_s, o_t, P_i)\|^2 \right), \tag{3.11}$$

where n is the number of 2D points (corners of the phantom) visible on the radiograph, p_i is the i^{th} 2D point, P_i is the corresponding 3D coordinate of the phantom model, prj is the 2D projection of a 3D point as defined in equations 3.1 and 3.2 (with the change that f and t_z defined in equation 3.2 should be

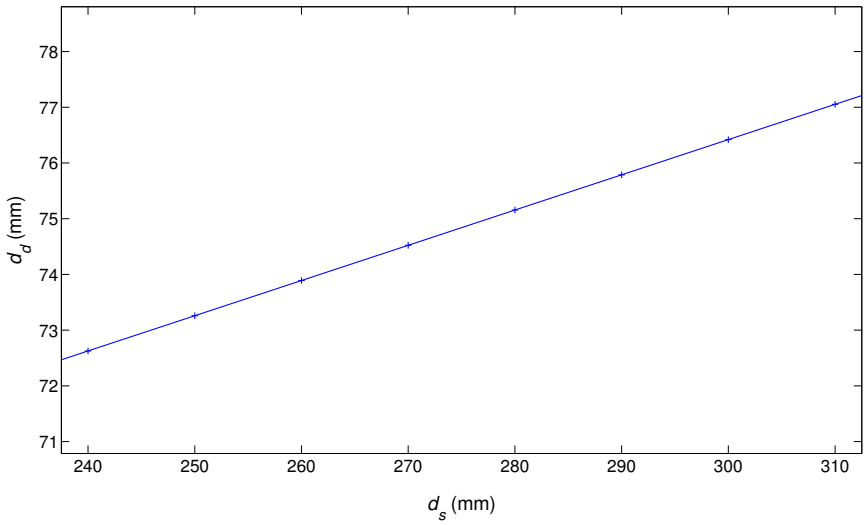


Figure 3.5: Relation between d_d and d_s for $d_m = 909$ mm on a real imaging system .

calculated using equations 3.8 and 3.9 respectively), and ζ^* is the optimised set of parameters.

The relation between d_d and d_s for a given d_m and for the kind of setup adopted is linear for the considered range of values (figure 3.5). The interception of the lines that describe this relation for a set of radiographs acquired at different distances d_m allows to find the system's values of d_d and d_s (figure 3.6). While two radiographs are enough for finding these values, taking several radiographs allows to increase robustness to noise on the input. When using more than two radiographs we have an overdetermined system where the intersection may not be perfect. Therefore, the system is solved in the linear least-squares sense.

3.4.2 Calibration parameters estimation

Having d_s and d_d calculated for a given system, converting the distance measured by the device (d_m) into f during an examination is very straightforward:

$$f = d_s + d_m + d_d. \tag{3.12}$$

In addition, the distance from the x-ray source to the table may be calculated as $d_s + d_m$, which may be used to initialize the value of parameter t_z (the distance from the x-ray source to the centre of the target being radiographed).

Regarding the remaining parameters, usually, the principal point (u_p, v_p) is located at the centre of radiographs (in pixels), the target structure is centred

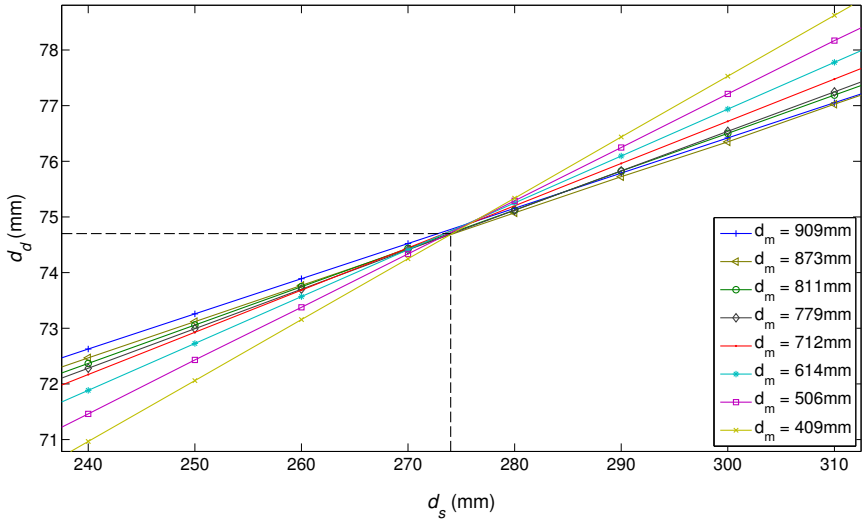


FIGURE 3.6: Interception of the relations between d_d and d_s for eight radiographs of the phantom acquired at different distances d_m on a real imaging system.

on the radiographs ($t_x \simeq 0$ mm, $t_y \simeq 0$ mm), and the orientation of the target structure is roughly known. For instance, in spine biplanar radiography, one may assume that in the frontal radiograph the subject is parallel to the x-ray detector (e.g. $\omega_{x_1} \simeq 0$, $\omega_{y_1} \simeq 0$, $\omega_{z_1} \simeq 0$, in radians), and in the lateral (s)he experiences a 90 degrees rotation around the Y axis (e.g. $\omega_{x_2} \simeq 0$, $\omega_{y_2} \simeq \pi/2$, $\omega_{z_2} \simeq 0$, in radians).

3.5 Algorithms

The next subsections describe the options made concerning algorithms selection for implementing the proposed method.

3.5.1 Corners detection

The 2D coordinates of the corners of the squares of the phantom on the radiographs were detected semi-automatically with the Camera Calibration Toolbox for Matlab (Bouguet, 2010). This toolbox provides a software tool that estimates the location of the corners of a chequerboard pattern independently of its orientation, given the 4 outer limits of the pattern as well as the number of squares per dimension. The 2D position of the corners was then optimised using the function

`cvFindCornerSubPix` of OpenCV (Bradski, 2000). This function locally refines the corners' location based on the gradient of the image.

3.5.2 Least-squares minimisation

A very common approach to solve the least-squares problem of minimising the retro-projection error is to use the Levenberg-Marquardt algorithm (Marquardt, 1963). Here, we propose using a trust-region-reflective algorithm for nonlinear least-squares problems (Coleman and Li, 1996). Trust-region optimisers have shown to be more computational efficient than the Levenberg-Marquardt algorithm in bundle adjustment problems Lourakis and Argyros (2005). Additionally, the adopted trust-region optimiser allows to define bounds for constraining the range of values of the parameters being optimised Coleman and Li (1996), which limit the search space of solutions. This is especially advantageous if the range of error of the calibration parameters is known, which is particularly true for the focal distance where the error should depend of the accuracy and precision of the measuring device.

The trust-region optimiser starts by defining a region around the initial solution. This region is approximated by a quadratic surface, for which a minimum can be directly computed, resulting on a new candidate solution. The algorithm then verifies if there is an actual improvement of the cost by evaluating the cost function with the candidate solution. If there is, the iteration is successful and, thus, the new solution is adopted and the size of the trust-region is increased for the next iteration; otherwise, the iteration is unsuccessful and, consequently, the size of the trust-region is decreased and the solution is not updated. These steps are repeated until convergence.

3.5.3 Triangulation of point-matches

Triangulation is used both when minimising the retro-projection error and when calculating the 3D coordinates of the final reconstruction from the 2D point-matches. The linear least-squares algorithm (Hartley and Sturm, 1997) was used for these purposes. This algorithm shows superior accuracy and robustness to noise than the commonly used mid-point algorithm (e.g. Kadoury et al. (2007b)) while remaining equally inexpensive to compute (Hartley and Sturm, 1997).

3.6 Method evaluation

This section describes the method assessment in radiological environment using a phantom with known geometry, and includes the description of the materials that were utilised.

3.6.1 Materials

Phantom

A phantom of stainless steel (AISI 304) was built with dimensions 380×380×1 mm^3 and laser cut squares of 20.0 mm, also spaced by 20.0 mm. The precision of the cutting was of ±0.2 mm. The coordinates of the corners of the squares lying on the first row and on the first column were measured with a Coordinate Measuring Machine (CMM) (±0.001 mm) for updating the 3D model of the phantom. A full measuring of the grid was not possible due to limitations of the available machinery.

Distance measuring device

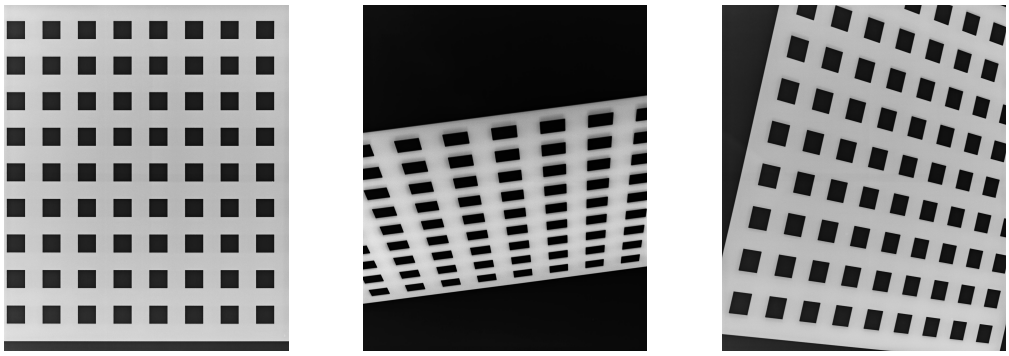
A laser rangefinder, Bosh DLE 50, Germany, was used on all experiments. This rangefinder has a typical error of ±1.5 mm, maximum error of ±3.0 mm, and range of operation of 0.05–50 m. Especial concern was taken to use a commercial, easy available, and affordable device (price ∼150 euros, Portugal, 2007).

Radiographs

Radiographs were acquired on a standard imaging system of Computer Radiography, Philips, Netherlands, at Serviço Médico de Imagem Computorizada, Porto, Portugal. Film size was 14″×17″ (355.6 × 431.8 mm), scanned with a sampling pitch of 175.0 μm/pixel (on both axes of the image), resulting in images with a resolution of 2010 × 2446 pixels.

3.6.2 Pre-calibration procedure

For determining parameters d_s and d_d, eight radiographs of the phantom were acquired with the x-ray source positioned at different distances from the table



(0, 0, 1183)
(0°, 0°, 0°)

(0, 0, 993)
(60°, 0°, 0°)

(0, 0, 993)
(0°, −40°, 10°)

Figure 3.7: Example of three radiographs of the phantom and the initial guess for translations in mm (above) and orientation in degrees (below).

(d_m ranging 409–909 mm) while keeping all other parameters constant. For each radiograph the distance between the x-ray device and the table (d_m) was measured using the laser rangefinder. All the corners visible in the radiographs were extracted semi-automatically and were used for calculating the projection error. The minimisation of the projection error was done using a trust-region-reflective algorithm (Coleman and Li, 1996). The precision of the pre-calibration procedure was evaluated by calculating the RMS of the residuals at the intersection point.

3.6.3 Calibration assessment

The same phantom from the previous procedure was utilised for evaluating the calibration method. This time, eight radiographs were taken with the phantom at different positions and with different orientations (figure 3.7). Distance d_m was the same for all setups ($d_m = 909$ mm). The eight radiographs were combined in a total of 17 pairs (out of 28 possible combinations). Only pairs of radiographs with considerably different pose were considered. Pairs of radiographs with near pose were discarded because are less tolerant to triangulation errors (when pose is similar triangulation lines tend to intersect at infinity because they are close to parallel).

For each setup, the phantom orientation was measured with a protractor and provided as initial guess to the optimiser in a 10° resolution scale. Parameter t_z was initialised with $(d_s + d_m)$ or $(d_s + d_m - 0.5 \times o_w)$, where o_w is the phantom width. The guess that was closer to the real value of t_z was chosen. Parameters t_x and t_y were always set to 0 mm. For better describing the input error, the

errors of the initial guess were estimated by comparison with the parameters that achieved optimal solutions when projecting a 3D model of the phantom onto the correspondent radiograph.

The bounds of the parameters for the trust-region optimiser were defined as ±100 mm for translation, ±10° for rotation, ±30 pixels for the principal point, and ±1.5 mm for the focal distance, in relation to the initial guess. The range of the focal distance was based on the specifications of the rangefinder.

Accuracy of the method was measured as the accuracy recovering the shape of the phantom. Therefore, accuracy is defined as the 3D Euclidian distances between the reconstructed corners of phantom and its 3D model (measured with a CMM) after aligning the reconstruction with the model. The rigid aligning is achieved using Horn's approach, which gives a closed-form solution for the case when the pairing between structures being aligned is known (Horn et al., 1987). This algorithm finds the translation and rotation that minimises the least-squares errors between the two sets and was shown to provide both stable and accurate rigid alignments in the presence of noise (Eggert et al., 1997). In fact, the study performed by Eggert et al. (1997) shown that under realistic conditions the main algorithms used for this task have similar performance.

Scale recovery

The accuracy of the method was evaluated with and without scale correction. For correcting scale, the scaling factor was calculated as the ratio between the real distance between two points of the phantom and the distance between the reconstructed 3D coordinates of the same points. Scaling is performed before aligning the reconstruction with the reference 3D model. When calculating errors for scaled reconstructions, 50 pairs of points were used for providing 50 scaling factors uniformly distributed by the part of the phantom that was visible in both radiographs. All reference distances had ∼40 mm, which is equivalent to two consecutive squares of the phantom. One reconstruction was performed per scaling factor for eliminating the effect of the location of the pair of points on the final result.

Ideal point-matches

An experiment was made to determine the method's performance when using accurately identified point-matches. This was done by using all the points of the phantom that were simultaneously visible on the two radiographs. The points were semi-automatically extracted as described in section 3.5.1.

Non-ideal point-matches

Another set of experiments was carried out for testing the calibration method in less optimal conditions concerning point-matches identification. This was done by adding uniformly distributed noise to the 2D coordinates of the phantom's corners that were extracted semi-automatically. Then, the previous experiment was repeated starting with no noise, and adding up to ± 15 pixels to each coordinate of every point. Additionally, the performance of the method with smaller sets of points was also evaluated.

Optimisers comparison and effect of constraining focal distance

A final experiment was done for comparing the performance of the proposed optimiser (Coleman and Li, 1996) against the commonly used Levenberg-Marquardt algorithm (Marquardt, 1963). Additionally, it was also tested the effect of relaxing the boundaries of the focal distance on the optimisation process to determine if using a rangefinder actually contributes to more accurate calibrations.

3.7 Results

3.7.1 Pre-calibration procedure

Figure 3.6 shows the relation between d_d and d_s for 8 radiographs with different d_m. At the intersection point, the RMS of the residuals for d_d was 0.03 mm and for d_s was 0.44 mm. RMS projection errors for the 8 radiographs using the values founded by the pre-calibration procedure were of 0.0004 pixels.

3.7.2 Errors of the initial guess of the calibration parameters

Focal length was calculated to be $f = 1257.7$ mm and the distance between the x-ray source and the table equal to 1183.0 mm. No ground truth was available for these two parameters as well as for the principal point. Average absolute orientation and translation errors were 1.7° and 20.2 mm respectively (table 3.1).

TABLE 3.1: Translation and orientation absolute errors of the initial guess.

	t_x (mm)	t_y (mm)	t_z (mm)	ω_x (deg)	ω_y (deg)	ω_z (deg)
Mean	17.7	21.0	21.9	0.9	1.7	2.4
SD	7.2	13.4	21.6	1.6	1.6	2.5
Max	30.5	43.2	58.6	3.8	3.9	6.9

3.7.3 Reconstruction accuracy with ideal point-matches

Mean 3D reconstruction errors were of 5.19 mm when not correcting scale, and 0.31 mm when correcting scale with a reference distance of ∼40 mm. This shows that, for the kind of setups that were tested, the method needs to know the distance between a pair of points in real world units for correcting scale. Therefore, the results of the following experiments include correction of scale with a reference distance. When correcting scale, 99% of the errors were inferior to 0.85 mm (figure 3.8).

3.7.4 Reconstruction accuracy with non-ideal point-matches

The relation between noise on the 2D point-matches and the 3D reconstruction error is close to linear ($R^2 = 0.993$) for the tested range of noise (figure 3.9). Additionally, the increase rate of the 3D reconstruction error is inferior to the rate on the input noise. For an uniformly distributed noise of 5 pixels the RMS 3D reconstruction error still remains inferior to 1 mm.

Decreasing the average number of point-matches from 199 to 33 produces an increase of $\sim 5\%$ on the 3D reconstruction error independently of the noise on the point-matches (table 3.2). When decreasing from 199 to 23 point-matches, the error varied from 9% to 19%, showing that for this amount of input the method is more sensible to noise on the point-matches.

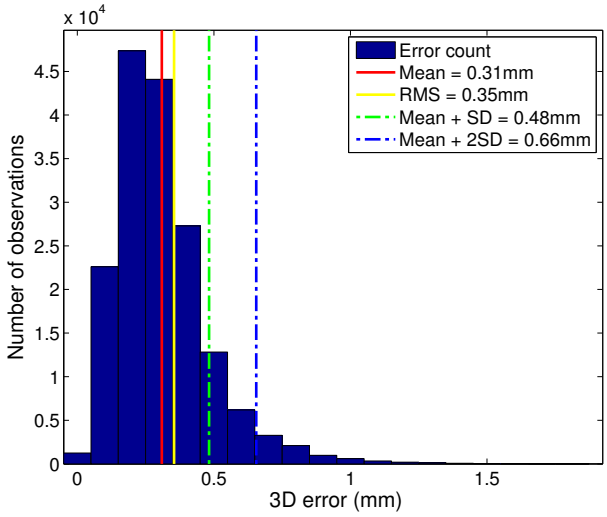


Figure 3.8: Histogram of 3D reconstruction errors for 17 pairs of radiographs using the maximum number of available point-matches.

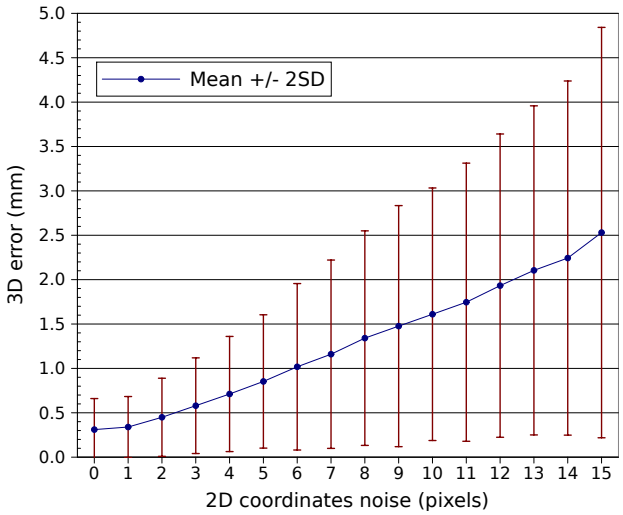


Figure 3.9: 3D reconstruction error vs. noise in point-matches location.

Table 3.2: 3D reconstruction error vs. number of point-matches and noise in point-matches location.

Number of matches	3D RMS reconstruction errors (mm)		
	No noise	±5 pixels noise	±10 pixels noise
199	0.36	0.93	1.76
67	0.36	0.98	1.80
33	0.38	0.98	1.86
23	0.39	1.21	2.10

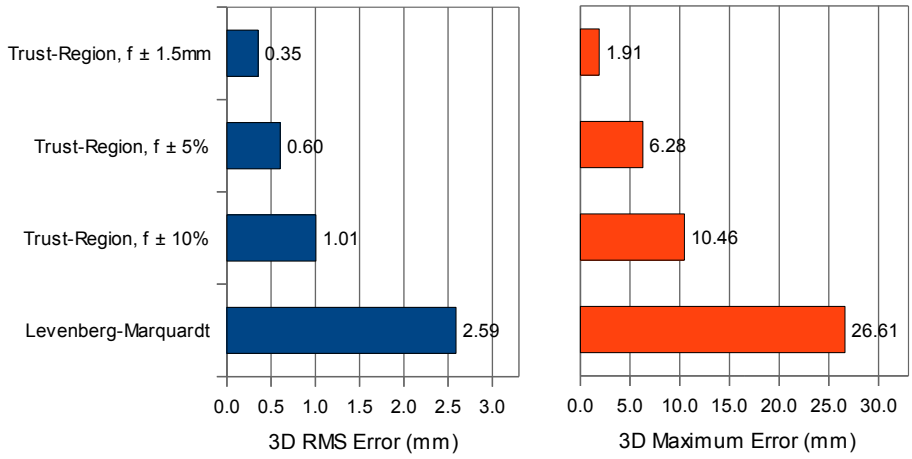


FIGURE 3.10: Comparison of optimisers and effect of the range of the bounds for the focal distance (f) on the trust-region optimiser.

3.7.5 Optimisers comparison and effect of constraining focal distance

Calibrations with the trust-region optimiser enabled to achieve 3D RMS reconstruction errors more than 7 times lower than when using the Levenberg-Marquardt optimiser (figure 3.10). While the trust-region optimiser was always able to converge to solutions where the reconstruction error was near the average error, the Levenberg-Marquardt optimiser was trapped in local minima for several setups that resulted in inaccurate reconstructions (figure 3.11).

When using the trust-region optimiser, relaxing the boundaries of the focal distance results on higher reconstruction errors (figure 3.12). Results show that relaxing the boundaries to $\pm 5\%$ of the initial value results in an increment of 1.7 times on the 3D RMS reconstruction error and of 3.3 times on the maximum error (figure 3.10). The higher the search range for the focal distance, the higher the reconstruction error. In addition, the maximum error when relaxing the focal distance by $\pm 5\%$ and $\pm 10\%$ is $\sim$10 times the RMS error, decreasing to $\sim$5 times when using the rangefinder specifications.

3.8 Discussion

Experiments show that the proposed method is able to accurately recover the 3D shape of a planar phantom from two radiographs showing different views of the phantom. However, in order to recover scale, a reference measure is needed

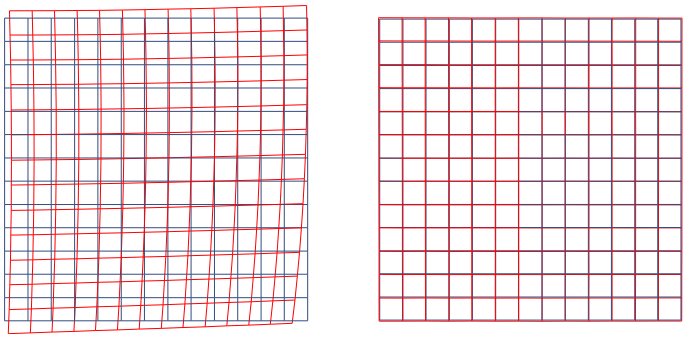
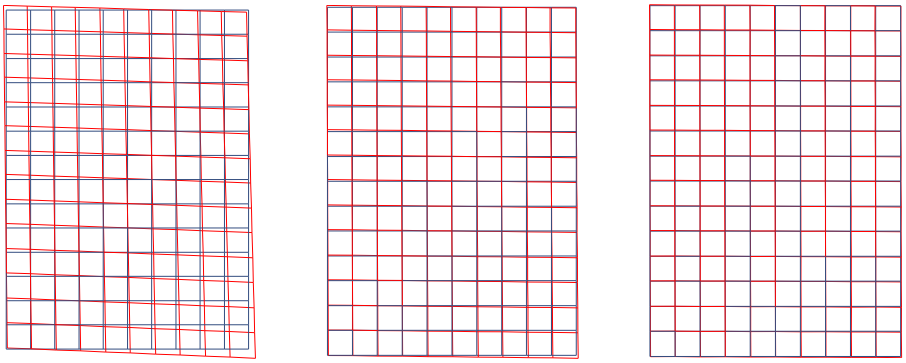


Figure 3.11: Reconstruction of a case where the Levenberg-Marquardt algorithm failed to capture the scene geometry (left) and the result with the trust-region algorithm for exactly the same conditions (right).



$f \pm 10\%$ $f \pm 5\%$ $f \pm 1.5$mm ($\sim$0.1%)

Figure 3.12: Illustration of the effect of the range of the bounds for the focal distance (f) when using the trust-region optimiser. The conditions for all reconstructions were exactly the same and $f = 1257.7$mm.

for finding the ratio between the size of the 3D reconstruction and the real size of the object. Even though, this still remains a low requirement when compared to other biplanar calibration methods and, due to limitations of the study highlighted at the end of this discussion, the hypothesis of recovering the scale without calibration objects on biplanar radiography of the spine is not discarded by the results of this chapter.

The sub-millimetric accuracy achieved by the proposed method was possible due to the calculation of a good initial guess using a distance measuring device coupled with the selection of an optimisation algorithm that enables to limit the search space of solutions, in particular for the focal distance. The region-based optimiser was shown to be superior to the commonly used Levenberg-Marquardt and restricting the focal distance was shown to improve reconstruction results, even when the initial guess of the focal length and t_z is calculated using the

proposed method. Experiments with erroneous focal lengths were not performed since varying the search range for this parameter was enough to show its sensibility.

Reconstruction errors are dependent of the errors identifying the point-matches; however, this relation is close to linear for the tested range of noise. The method has shown robustness when decreasing the number of point-matches by a factor of 6, with an increase of the error of only ∼5% independently of identification errors. Decreasing the number of matches even further is possible at the cost of increasing reconstruction errors, especially when the point-matches are noisy.

Results presented here also compare well with the reconstruction of steel pellets presented in (Pearcy, 1985) with a fixed biplanar setup composed by two x-ray sources and two detectors with calibration based on DLT. In particular, the method proposed here achieved an RMS error of 0.35 mmm while the fixed system presented an RMS error of 0.44 mm.

The results found here may not be generalized to spine radiography. The phantom object does not resemble a spine, x-ray detectors in spine radiography are typically higher for capturing the complete spine, and focal distance can also take higher values. Additionally, since the phantom is planar it was not possible to place it in a lateral pose, which is typical of biplanar radiography of the spine. Moreover, the phantom normally provides point-matches spread all over the radiographs, which does not happen with the spine. Thus, results presented in this chapter cannot be compared with calibration studies of the spine. Nevertheless, results are promising and valid for general biplanar radiography.

3.9 Summary

This chapter presented a novel method for the calibration of biplanar radiographs that is based on the estimation of an initial guess of the calibration parameters using a distance that can be easily measured on site, followed by a minimisation of the retro-projection error of a set of point-matches with a trust-region optimiser. A pre-calibration procedure that enables using a distance measuring device was also proposed. Experimental results with a radiopaque phantom on radiological environment shown that the method requires to know a reference distance in order to properly recover scale. Under this condition the method achieved sub-millimetric accuracy and shows robustness to noise on the

point-matches even when just a fraction of the points that compose the phantom are used to optimise the calibration parameters. Results also highlighted the importance of the precision of the measuring device as well as the importance of choosing an adequate optimisation algorithm.

CHAPTER 4

Calibration of biplanar radiographs of the spine

In the previous chapter a novel calibration method for general biplanar radiography was presented. This chapter describes adaptations and studies of the method for the particular case of spine radiography (Moura et al., 2010a). Additionally, several issues like the positioning and size of the calibration object, the precision of the distance measuring device, and the impact of the steps of the calibration method are addressed. The method is validated both using simulations with 3D *in vivo* data, as well as *in vitro* experiments with a dried spine. Finally, results are discussed and compared with related work.

4.1 Calibration system for 3D reconstruction of the spine

In order to use the calibration method proposed in chapter 3 for spine radiography with real patients, several issues have to be addressed, such as patient positioning and radiographs acquisition protocol, among others. This section proposes adaptations to the calibration method and integrates it with a 3D reconstruction system of the spine (figure 4.1). The complete system will be described here and in section 4.2.2 the influence of each one of the calibration steps will be

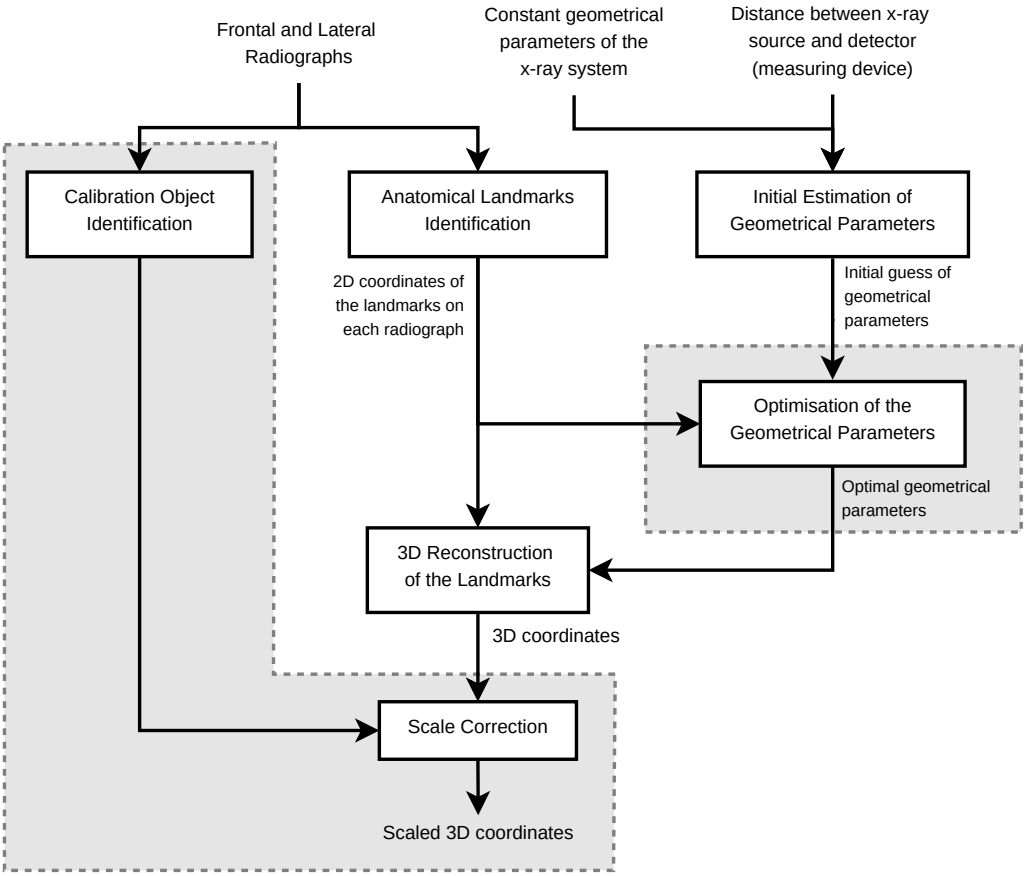


Figure 4.1: Data flow diagram of the 3D reconstruction system. The impact of the calibration processes that are inside gray boxes will be assessed on section 4.2.2.

assessed to determine their impact on the accuracy of 3D reconstructions of the spine.

The input data are two radiographs (Frontal and Lateral) of the subject's spine in digital format, the distance between the x-ray source and the x-ray table (d_m) measured with a measuring device during the radiographs acquisition, and a set of geometric parameters of the radiographic system (d_d and d_s) that are constant and therefore only need to be determined once for a given system with the pre-calibration procedure described in section 3.4.1. Parameters d_d and d_s together with the distance from the measuring device (d_m) are used to obtain an initial guess of the calibration parameters, which are then optimised by minimising the retro-projection error of a set of landmarks (section 3.3). Finally, scale is corrected using a simple calibration object of known geometry.

4.1.1 Radiographs acquisition procedure

The proposed method requires two orthogonal radiographs of the spine, one AP (Antero-posterior) or PA (Postero-anterior), and one Lateral (Left-right or Right-left). These planes are commonly used by physicians for the follow-up of spinal deformities, such as scoliosis (Cassar-Pullicino and Eisenstein, 2002; Greenspan, 2004), and, therefore, the proposed method does not subject patients to additional radiation.

The radiographs acquisition procedure starts by positioning the subject. For guaranteeing a proper positioning without using a rotatory platform, we suggest the same option as Kadoury et al. (2007b), which consists on using markers on the floor for subjects to place their feet (figure 4.2). These markers help stabilising subjects by making legs to be apart by a considerable distance. Additionally, they allow to have an approximate value of the distance between the subject's spine and the x-ray table (d_p). After positioning the subject, the technician should adjust the distance between the x-ray source and the subject for best fitting the region of interest in the radiographs. After selecting this distance, the first radiograph takes place, followed by a 90 degrees rotation of the subject with the help of the foot-markers for acquiring the second radiograph. Finally, the subject leaves the system, which allows the technician to measure the distance between x-ray source and the x-ray table (d_m). We propose measuring this distance with a laser rangefinder. This kind of devices allow to quickly and easily perform measurements with millimetric accuracy, while being inexpensive.

4.1.2 Point-matches identification

The calibration method needs a set of point-matches on the two radiographs for optimising the calibration parameters. Here, it is assumed that the 2D co-ordinates of a set of anatomical landmarks are available for each radiograph. These landmarks are six corresponding points per vertebrae that are visible on both radiographs (i.e. centre of superior and inferior endplates, and the superior and inferior extremities of the pedicles) and that are widely used on methods for reconstructing the spine from biplanar radiography (e.g. Aubin et al. (1995); Mitton et al. (2000); Kadoury et al. (2007a,b); Delorme et al. (2003); Mitulescu et al. (2002); Bras et al. (2003)). Such landmarks may be manually identified or a less supervised method may be used that does not rely on the calibration

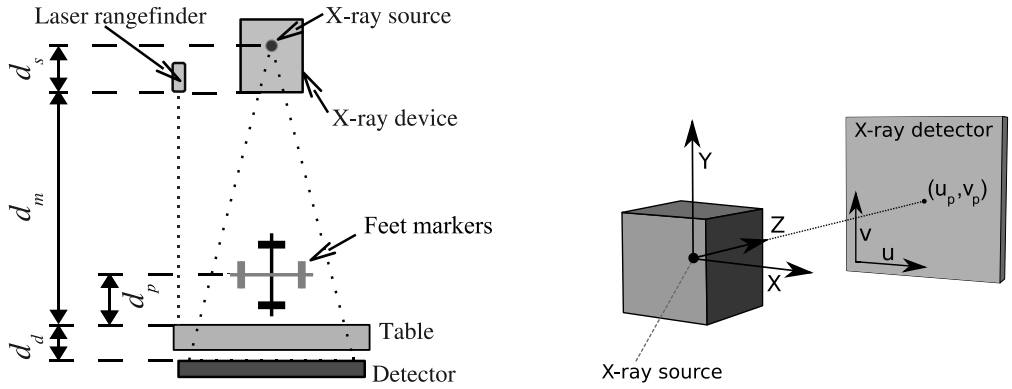


Figure 4.2: On the left, an illustration of a conventional radiographic imaging system (top view) with a laser rangefinder attached. This figure is similar to figure 3.3 but here the markers where subjects should place their feet are also represented (gray markers for the frontal radiograph; black markers for the lateral radiograph), as well as distance d_p. The table is in upright position. The figure on the right shows the coordinate system of reference; the 3D coordinate system is located at the centre of the x-ray source with the Y axis in the caudal-cranial direction.

matrices to obtain the 2D coordinates of these landmarks (e.g. Vaiton et al. (2004)).

The six points per vertebrae will also be used for computing the 3D reconstruction of the spine by calculating their 3D coordinates using the linear least-squares algorithm (Hartley and Sturm, 1997).

4.1.3 Initial guess of the calibration parameters

The assumptions made regarding parameters initialization for general biplanar radiography (section 3.4.2) are valid for the case of the spine, i.e. one may assume that the spine is roughly centred on the radiographs and that the source is roughly pointing to centre of the x-ray detector. Regarding the spine orientation on the two radiographs, let the coordinate system of reference be located at the x-ray source as defined in figure 4.2, one may assume that in the frontal radiograph the subject is parallel to the x-ray detector (e.g. $\omega_{x_1} \simeq 0$, $\omega_{y_1} \simeq 0$, $\omega_{z_1} \simeq 0$, in radians), and in the lateral (s)he experiences a 90 degrees rotation around the Y axis (e.g. $\omega_{x_2} \simeq 0$, $\omega_{y_2} \simeq \pi/2$, $\omega_{z_2} \simeq 0$, in radians).

Finally, the initialization of the parameters that depend of the distance measured by the device (d_m) is done in the same way as before, with the exception that the feet markers may now be used for estimating the distance between the

subject and the x-ray source (t_{z_i}). Therefore, for each radiograph i:

$$f_i = d_s + d_m + d_d, \qquad t_{z_i} = d_s + d_m - d_p, \qquad \text{for } i = 1, 2. \tag{4.1}$$

While the focal distance is accurately calculated, t_{z_i} is only an approximated value that depends of the subject's positioning and anatomy.

4.1.4 Scale correction

Experimental results shown that the method was not able of handling scale when reconstructing a planar phantom, but that only a reference distance is needed for properly scaling reconstructions (section 3.7.3). Therefore, for guaranteeing that scale is properly recovered the subject wears a calibration object that undergoes the same geometrical transformation that the subject experiences when rotating from the first to the second radiograph. For minimal impact on the content of radiographs, the object only needs to have two radiopaque pellets at a known distance. Additionally, it should be placed on the lumbar area of the backs of the subject for minimal overlapping with bone structures (i.e. rib cage) and facilitating its identification. The scaling factor may be calculated as the ratio between the real distance between the two pellets and the distance between the reconstructed 3D coordinates of the pellets identified on both radiographs. The impact of the size and position of the calibration object will be addressed in section 4.2.1.

4.2 Simulation studies with in vivo 3D data

Computer simulation was employed for assessing the method by simulating radiographic exams using data from an *in vivo* CT scan. The CT scan was of a woman with 77 years old * and captures the complete thoracic and lumbar spine with voxel size of $0.4 \times 0.4 \times 0.5$ mm^3. The six anatomical landmarks per vertebrae were manually identified by a human expert on the CT scan for vertebrae T1 to L5, and constitute the ground truth for this study. For simulating radiographs acquisitions on realistic conditions, Gaussian noise was added to the geometric parameters and to the landmarks. First, noise was added to the reference 3D landmarks and to the geometric parameters of the frontal view and the noisy

∗. This exam was retrieved from a database of clinical exams and, therefore, it was not especially acquired for this experiment nor for clinical trials.

landmarks were projected with the noisy parameters for simulating the input (2D points) of this view. Then, the same was done to the lateral view, that is, noise was added to the geometric parameters of the lateral view and again to the reference 3D landmarks for simulating both errors on establishing the stereo-correspondence of the point-matches as well as identifying their real location. No noise was added to the landmarks of the calibration object since radiopaque pellets can be precisely identified on radiographs. Unless specified otherwise, the projection planes were the Antero-posterior and Left-right. Figure 4.3 illustrates a simulation of a single exam and the correspondent 3D reconstruction.

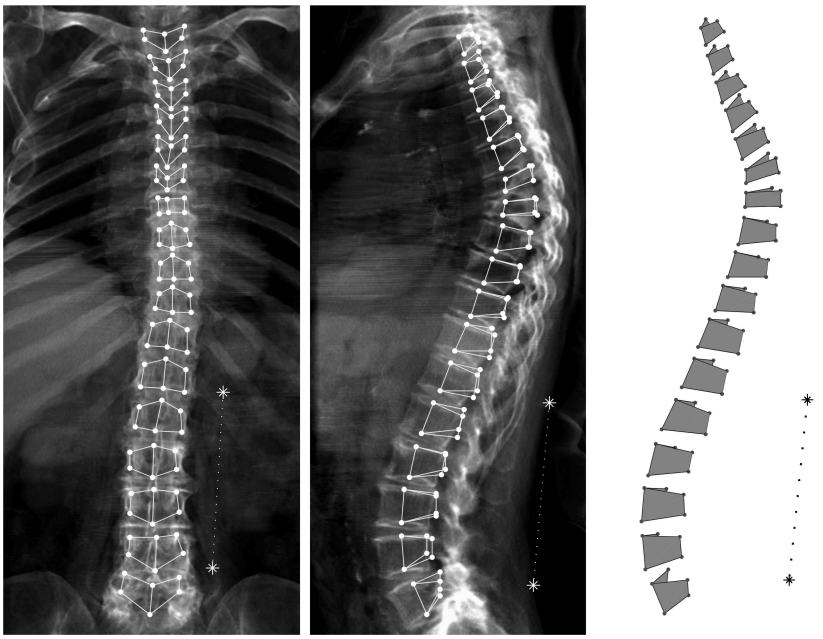


Figure 4.3: On the left, an illustration of the input of a simulation: the anatomical landmarks (filled dots) and the landmarks of the calibration object (asterisks); digitally reconstructed radiographs were rendered from the CT scan for illustration purposes. On the right, a perspective view of the 3D reconstruction of the input landmarks after calibration.

Two levels of noise were used in the simulations (table 4.1). In the first (Controlled Setup), the same amount of noise used in Kadoury et al. (2007b) was considered for all geometric parameters except focal length, which can be accurately calculated with the proposed method using a rangefinder. Therefore, the noise on the focal length was based on the accuracy specifications of an off-the-shelf laser rangefinder (S.D. of the error of 1.5 mm). In a second experiment (Pessimistic Scenario) the geometric errors were increased (with the exception of focal length) using the maximum expected variation on this kind of clinical

TABLE 4.1: Standard deviation of the Gaussian noise that was added to the geometric parameters on simulations.

Geometric parameters	Noise 1^{st} experience (Controlled Setup)	Noise 2^{nd} experience (Pessimistic Scenario)
Rotation $(\omega_x, \omega_y, \omega_z)$	0.5°	2.0°
Translation (t_x, t_y, t_z)	4.0 mm	8.0 mm
Principal point (u_p, v_p)	2.0 mm	4.0 mm
Focal distance (f)	1.5 mm	1.5 mm

setups (Kadoury et al., 2007b). For both scenarios the amount of noise added to the anatomical landmarks had standard deviation of 1 mm, which is the expected for landmarks identified by a human expert (Kadoury et al., 2007b). The bounds used on the optimisation algorithm were set to 4 times the S.D. of the noise added to the corresponding geometric parameters, with the exception of the focal length that was not optimised.

For each of the following experiments, 100 trials were simulated and, for each trial, random Gaussian noise was added to both parameters and landmarks. Two measures were used for assessing the reconstruction quality: the 3D reconstruction error of each landmark (after rigidly aligning the reconstruction with the ground truth), and the 3D spinal length (Papin et al., 1999). This last clinical index was included on this study because of difficulties of previous methods (Kadoury et al., 2007a) on determining it. It is calculated by summing the Euclidian distances between every pair of consecutive vertebral bodies' centres, which, on their turn, are calculated as the midpoint between the superior and inferior centres of the vertebra's endplates.

4.2.1 Positioning and size of the calibration object

A set of simulations were performed to determine the effect on the reconstructions of location, orientation and size of the calibration object. All experiments were done with the noise settings of the Pessimistic Scenario (table 4.1). First, the impact of the object's orientation was assessed by measuring the RMS 3D reconstruction error obtained by simulating a bar with 120 mm length placed at different orientations on the lumbar area of the subject's backs. Inclinations on the frontal and lateral planes were tested independently with the inclination of the bar ranging from 0° to 180° with intervals of 15°. For eliminating the effect of the location of the calibration object, for each of the 100 trials, the object was

placed at 27 different locations uniformly distributed over the subject's lumbar area.

After finding the optimal object orientation, a second experiment was done for determining the effect of the position of the calibration object on the reconstruction error. Again, for each trial, the object was placed at 27 different locations uniformly distributed over the subject's lumbar area but, this time, errors were calculated independently for each position. The object was always placed with optimal orientation. Additionally, all combinations of frontal and lateral planes, i.e. antero-posterior and postero-anterior with right-left and left-right, were experimented for determining if the subject orientation had any effect on the calibration object positioning.

Finally, the impact of the size of the calibration object was assessed by running simulations with different sizes of the calibration object, ranging from 20 mm to 160 mm in increments of 20 mm. For each trial, the calibration object was placed randomly within the optimal area and with optimal orientation.

4.2.2 Impact of the calibration components

For determining the gain and need of using the different components of the proposed method (i.e. the rangefinder, the optimisation process, and the calibration object), the following setups were tested:

1. **Initial Error**: the reconstruction was performed by triangulating the noisy anatomical landmarks, with no optimisation of the geometric parameters and no scale correction. For simulating that no rangefinder was being used, the noise on focal distance was set to S.D. 20 mm on the first experiment (Controlled Setup) and to 40 mm on the second experiment (Pessimistic Scenario) (Kadoury et al., 2007b);
2. **Rangefinder**: same as Setup 1, but the S.D. of the focal distance was decreased to 1.5 mm to simulate the use of a simple rangefinder;
3. **Rangefinder and Calibration Object**: same as Setup 2, but after triangulating the landmarks they were scaled using the calibration object;
4. **Rangefinder and Optimisation**: same as Setup 2, but before triangulating the landmarks the geometric parameters were optimised;
5. **Complete Process**: equivalent to Setup 4, followed by scale correction with the calibration object.

For setups 3 and 5 a calibration object with 120 mm length was simulated. For each trial, it was randomly placed within the optimal area of the subject's back. For simulating real conditions, Gaussian noise was also added to the object optimal orientation (standard deviation of 2.4°).

Finally, the accuracy of the method recovering the absolute orientation of the 3D reconstruction was evaluated by recording the rigid transform that provided the best alignment between the 3D reconstruction and the ground truth. The transform was calculated using Horn's algorithm (Horn et al., 1987).

4.2.3 Effect of the precision of the measuring device

An experiment was conducted to evaluate the effect of the precision of the measuring device on the quality of the reconstruction. This experiment was done using the noise settings of the Pessimistic Scenario, simulating an increasing range of noise on the measuring device, ranging from Gaussian noise with standard deviation of 0 mm to 100 mm. All the conditions of the previous experiment were replicated, including the different setups of the reconstruction system.

4.3 In vitro validation

The method was validated using 17 dried vertebrae (from T1 to L5) of a human with unknown age, which were disposed with a holder in order to resemble a typical spine (figure 4.4). A calibration bar of length 122.3 mm was also placed in the holder for scaling the final 3D reconstruction. The spine was first scanned using CT (voxel size of $0.4\times0.4\times0.3$ mm^3) and then radiographed in the AP and Right-left planes (pixel size 0.1750×0.1750 mm^2) using a film holder for large radiographs. The holder enables to have 3 films with dimensions of $14''\times14''$, which are then scanned separately and finally merged. For this experiment only 2 films were needed to capture the complete spine. When acquiring the radiographs, an off-the-shelf laser rangefinder (a Bosh DLE 50, Netherlands, which has a typical error of ±1.5 mm, maximum error of ±3.0 mm, and range of operation of 0.05–50m) was used to measure distance d_m (figure 4.2). Then, all landmarks were identified in both CT and radiographs by an expert. The proposed method was used on the biplanar radiographic data to determine the calibration parameters and to compute the 3D coordinates of the anatomical landmarks, which were then rigidly aligned and compared with the landmarks of the CT scan. The 3D

spinal length was also evaluated. In addition, the same setups of the method that were used in the simulations (section 4.2.2) were included in this experiment for confronting both studies.

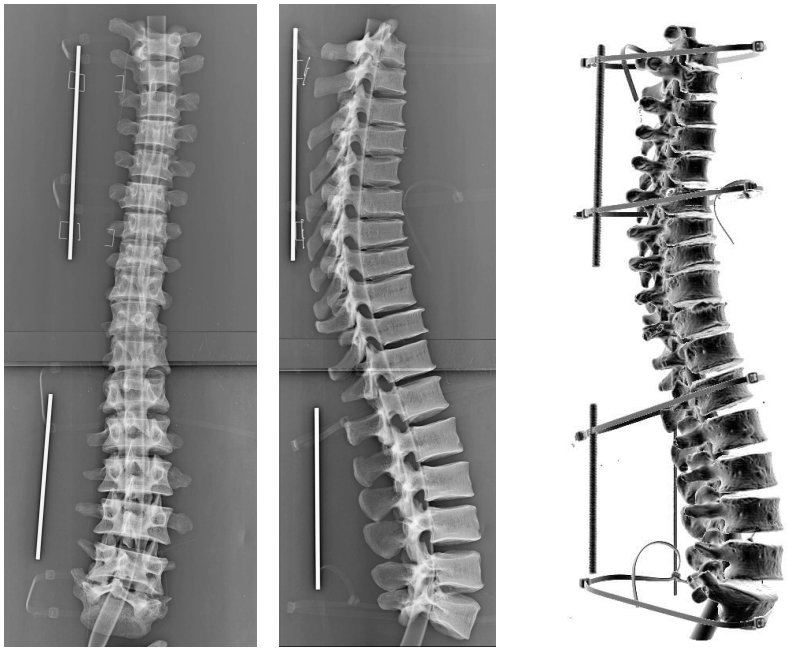


Figure 4.4: On the left, radiographs of the dried vertebrae with the calibration bar (AP and Lateral) and, on the right, reconstruction from the CT scan. Only the bottom bar was used on the experiments reported here.

4.4 Results

4.4.1 Positioning and size of the calibration object

Simulation results show that the proposed method achieves the lowest errors when the calibration bar is positioned approximately with vertical orientation (figure 4.5). Therefore, in the next simulations the object was placed vertically. Figure 4.6 shows that the position of the calibration object also affects the 3D reconstruction. For the tested range of positions, the 3D RMS reconstruction error varied 0.1 mm and the error on the spinal length 1.2 mm. This variation is particularly observable along the medial-lateral axis of the subject. Experimental results with different combinations of lateral and frontal radiographs show that the position where the lower errors are found depends of the orientation of the subject on the two acquisitions (table 4.2).

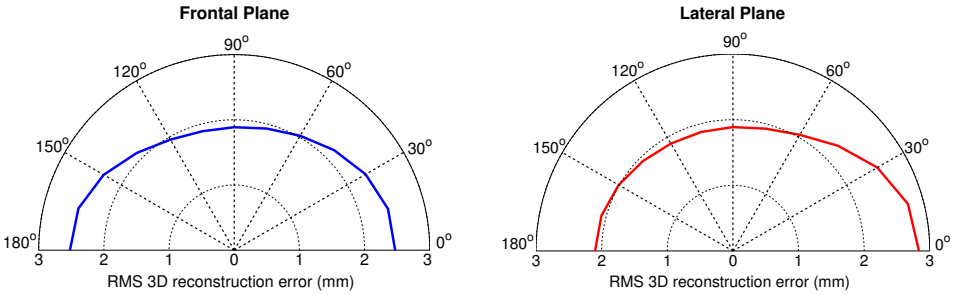


FIGURE 4.5: Effect of the orientation of the calibration object on the 3D reconstruction error. An inclination of 90° means that the calibration object is vertically oriented, while 0° means that it is horizontally oriented.

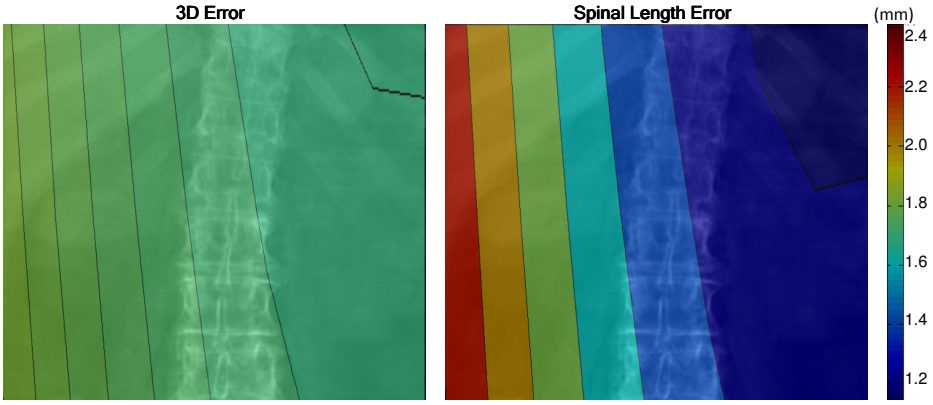


FIGURE 4.6: Illustration of the effect of the position of the calibration object on the 3D reconstruction error. The colored areas show the reconstruction errors that are obtained when the calibration object is placed inside them, with hotter colors meaning higher errors, according to the color bar at the right.

TABLE 4.2: Calibration object position, in terms of side of the subject, that shows the lowest reconstruction errors according to the subject's orientation on the frontal and lateral radiographs.

	Left-right x-ray	Right-left x-ray
Antero-posterior x-ray	Left side	Right side
Postero-anterior x-ray	Right side	Left side

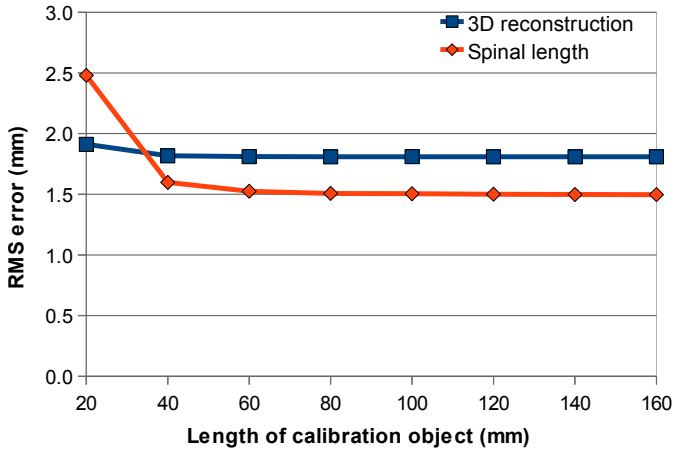


Figure 4.7: Effect of the size of the calibration object on the 3D reconstruction error.

Finally, simulations also show that the size of the calibration object affects the reconstruction accuracy (figure 4.7). In particular, the higher the length of the object, the lower the reconstruction errors, although a length higher than 80 mm does not produce observable improvements on any of the evaluated indices. Spinal length show higher improvements with the size of the calibration object than the 3D reconstruction error.

4.4.2 Method accuracy and impact of the calibration components

Results of the simulated experiments (figure 4.8) show RMS 3D reconstruction errors between 1.6 mm and 1.7 mm for the Controlled Setup and between 1.8 mm and 2.0 mm for the Pessimistic Scenario (noise range for each scenario is detailed on table 4.1). These results revealed that calculating the focal distance with the proposed method using the input of the rangefinder enables, by itself, to achieve comparable 3D reconstruction errors to the complete process Concerning the calculation of the spinal length, the use of a rangefinder *per se* decreased errors by a factor of 3.9 on the Controlled Setup and 4.6 in the Pessimistic Scenario. Correcting scale with the calibration object enabled to decrease spinal length errors from 2.7 mm to 2.5 mm, while the complete process decreased errors to 1.5 mm. Optimising the retro-projection error *per se* does not improve the results obtained when using the initial guess of the rangefinder directly, but when combined with scale correction provides the best results. Additionally, calibration decreases errors on the estimation of the absolute orientation of the reconstruction (table 4.3).

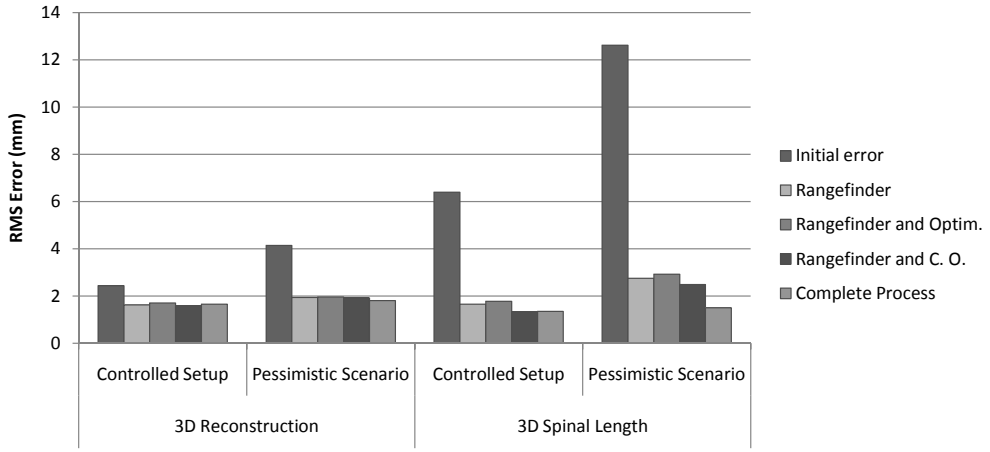


FIGURE 4.8: Simulation results for the comparison of 3D Reconstruction and Spinal Length errors of different configurations of the proposed method (described in section 4.2.2), for two different scenarios (detailed in table 4.1). (C.O. – Calibration Object; Optim. – Optimisation.)

TABLE 4.3: Simulations results for absolute orientation errors on the Pessimistic Scenario for different configurations of the proposed method (described in section 4.2.2).

	RMS orientation errors (degrees)			
Setup	X	Y	Z	Mean
Rangefinder	1.8	1.8	2.1	1.9
Rangefinder and Optimisation	1.6	1.4	1.8	1.6
Rangefinder and Calibration Object	1.6	1.9	1.6	1.7
Complete Process	1.4	1.8	1.5	1.6

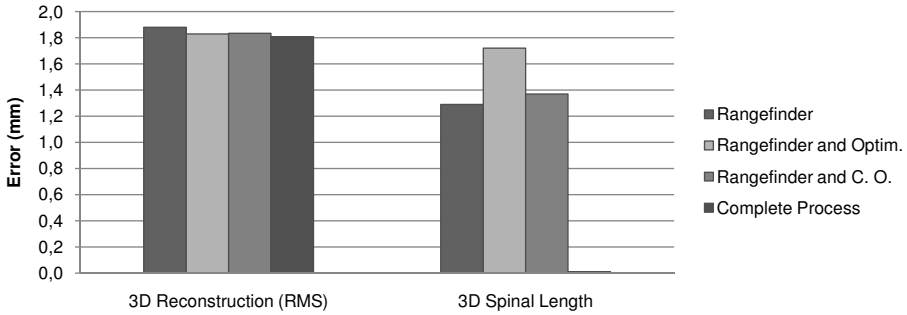


FIGURE 4.9: *In vitro* results for the comparison of 3D Reconstruction and Spinal Length errors of different configurations of the proposed method (described in section 4.2.2), for two different scenarios (detailed in table 4.1). The error of the spinal length when using the complete process is hardly visible because it only scored 0.01 mm. (C.O. – Calibration Object; Optim. – Optimisation.)

Table 4.4: 3D reconstruction errors (mean±S.D. in mm) of the *in vitro* validation for two configurations of the proposed method (described in section 4.2.2), and comparison with *in vitro* results using a large calibration apparatus (Aubin et al., 1997), and with a recent method based on high-level primitives that uses a small calibration object (Kadoury et al., 2010).

	Rangefinder	Complete Process	Large apparatus Aubin (1997)	High-level primitives Kadoury (2010)
Plates	1.7±0.6	1.7±0.6	1.5±0.7	1.9±0.9
Pedicles	1.8±0.8	1.6±0.8	1.2±0.7	2.9±1.3

In vitro results (figure 4.9) are consistent with simulation results, with 3D RMS reconstruction errors ranging from 1.8 mm to 1.9 mm. Additionally, the optimisation of the retro-projection error by itself did not improve results unless coupled with scale correction, like was observable on simulations. Once more the complete process obtained the best results. Differences on the 3D reconstruction errors among the different setups of the calibration method were found to be dominated by the pedicles landmarks (table 4.4).

4.4.3 Effect of the precision of the measuring device

Simulation results show that the 3D reconstruction error depends of the precision of the distance measuring device (figure 4.10). Using a calibration object provides superior robustness to noise on the input distance and the complete process always shows more accurate results. Regarding the spinal length, the calibration object again provides higher robustness and only with noise of 90 mm the complete process starts increasing the error calculating this index. Combining optimisation with scale correction improves the calculation of the spinal length on average 1 mm when compared with only using the calibration object. When not correcting scale, optimising the calibration parameters increases the spinal length error. Only using the initial guess provided by the proposed method without optimising parameters nor correcting scale shows stable results until approximately a precision of 2 mm of the measuring device. In this range, errors calculating the spinal length are on average 0.2 mm higher than when adding scale correction and 1.2 mm than when using the complete process. It is important to mention that on these experiments focal length was never optimised because using a search range proportional to the noise always increased both reconstruction and spinal length errors.

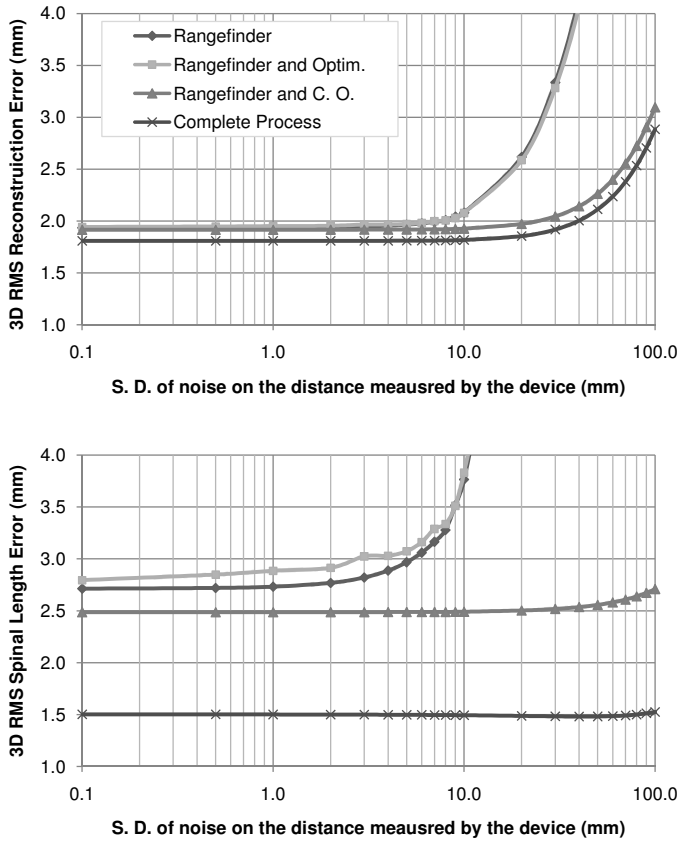


FIGURE 4.10: Simulation results for the evaluation of the effect of the precision of the distance measuring device on the 3D reconstruction error (top) and on the 3D spinal length (bottom), for different configuration of the proposed method (described in section 4.2.2). (C.O. – Calibration Object; Optim. – Optimisation.)

4.5 Discussion

Results clearly show that calculating focal length with the proposed method based on a distance easily measured on site (d_m) improves the recovery of scale of the reconstructions when not using calibration objects. In the experiments on the Pessimistic Scenario that simulated using only the initial guess of the calibration parameters for calculating the 3D reconstruction, the error calculating the spinal length was 2.75 mm. This is a major improvement over previous work, since results obtained by calibration methods that do not use calibration objects scored an error of 14.19 mm on this index (for *in vivo* experiments) (Kadoury et al., 2007a). Additionally, in all experiments, the 3D reconstruction error was comparable to the error after calibrating and scaling the reconstruction. However, for best results, simulations show that when using the initial guess directly,

the precision of the measuring device should not be inferior to 1 mm. Thus, the rangefinder that was used in the *in vitro* experiments is close to optimum while being affordable. Moreover, this validates the hypothesis that including an accurate estimation of distance d_m in the geometrical model is a key component for recovering scale and to eventually simplifying the method by discarding the scale correction and the optimisation of the parameters.

Simulations show that when using a calibration object for correcting scale, the quality of the reconstruction depends of the object's orientation, location and size. Placing the object vertically provides the lower 3D reconstruction errors. We believe that this is related to the fact that the points that represent the spine are mainly spread along the vertical (Y) axis. The Y axis is the one that shows higher variation and therefore roughly placing the calibration object along this axis should provide the better overall correction. With the object properly placed, its location also influences the reconstruction quality, in particular the spinal length. In fact, the 3D reconstruction error of the anatomical landmarks remains practically independent of the object location. This finding shows that the error calculating the spinal length may be doubled if the object is not properly positioned. Simulations also show that the subject's orientation influences the optimal positioning of the object, but numerical simulations allowed to determine that near optimal positioning may be achieved by simply placing the object in one of the sides of the subject's backs according to the acquired planes (PA vs AP, Left-right vs Righ-left). Finally, the size of the calibration object also influences reconstructions quality in particular the spinal length. A length of 80 mm is sufficient for achieving best results (RMS error calculating the spinal length of 1.5 mm for the Pessimistic Scenario), and higher lengths may be used without jeopardising results for reducing the impact of eventual identification errors of the object. Following these guidelines concerning the calibration object makes the spinal length calculation robust to noise on the distances provided by the measuring device. The 3D reconstruction error still remains dependent of the precision of the device, starting to deteriorate with noise higher than 10 mm. Even though, using a calibration object allows to use a distance measuring device with lower precision than the one used on the *in vitro* experiments.

From these experiments it is also possible to conclude that, when using a rangefinder of a precision similar to the one used on *in vitro* experiments, scaling the reconstruction using a small calibration object and optimising the geometrical parameters only seem to be manifestly advantageous when used together. This is especially observable on the simulation of the Pessimistic Scenario, where only

the complete version of the process achieved a considerably lower error on the spinal length when compared to the version that only used the initial guess of the parameters directly. Therefore, on further comparisons with other methods, only two variants of the method will be considered: a) using directly the initial guess calculated with the output of the rangefinder, and b) the complete process.

Reconstruction errors for the simulations on the Controlled Setup (1.6 mm when only using the initial guess and 1.7 mm for the complete process) were slightly lower than the 1.8 mm reconstruction error achieved on the simulations presented in Kadoury et al. (2007b) with similar noise levels (with the exception of focal length that has less noise in our case due to the use of a rangefinder). Despite this difference may be explained by factors that are difficult to replicate (e.g. source of 3D data), it shows that the methods have a comparable performance. However, while the method proposed in Kadoury et al. (2007b) needs a set of landmarks manually identified for optimising the geometrical parameters and uses a calibration object of 100 mm $\times$ 70 mm that slightly overlaps bone structures, it has been shown here that with only a measure from a rangefinder, similar performances are achieved in terms of 3D reconstruction error.

Simulation results also show that the error on the spinal length also compares well with the *in vivo* results of Kadoury et al. (2007b) where the authors obtained a RMS error of 2.05 mm. The complete version of the proposed method achieved a lower error on the Pessimistic Scenario (1.50 mm), with the additional advantage of using a calibration object with lower impact on the content of radiographs, and the disadvantage of needing a rangefinder. When using directly the initial guess the error of the spinal length was higher (2.75 mm). However, it is difficult to make a fair comparison with *in vivo* studies since there is no ground truth available. Additionally, these studies are performed with scoliotic spines while here we have used non-pathological spines. Therefore, further studies with *in vivo* data from scoliotic patients are required for a proper comparison.

Regarding robustness to the initial guess of the calibration parameters, simulations show that both variants handled an increase of noise by a factor of 2 in most of the parameters and by a factor of 4 on rotation (from the Controlled Setup to the Pessimistic Scenario) by only increasing the RMS 3D reconstruction error by factors of 1.09 on the complete version of the method, and 1.19 when using the initial guess directly. In terms of the spinal length, the method shown higher sensibility, especially when only using the initial guess calculated with the rangefinder (increasing factor of 1.11 vs 1.66). This shows, as expected,

that when using only a rangefinder, the method is more sensible to errors on the initial estimation of the geometric parameters and, therefore, proper subject positioning is more crucial than in the complete version of the proposed method.

In vitro results (table 4.4) also compare well with results from large calibration apparatus, like is the case of Aubin et al. (1997). As expected, the 3D reconstruction errors of the proposed method are higher (mean error of 1.7 mm vs 1.3 mm), since the method presented in Aubin et al. (1997) uses a calibration cage that completely surrounds the subject, and a rotatory platform that guarantees low rotation errors when positioning subjects. Nevertheless, the accuracy of the proposed method remains acceptable for clinical assessment. Additionally, the proposed method offers several advantages that may justify the loss of exactness, such as, much lower costs, superior user-friendliness, compatibility with standard radiological systems, and much less (or no) artifacts on radiographs overlapping anatomical structures. Additionally, while the complete version of the proposed method is able of handling motion of the subject between radiographs, the same does not happen with the cage method (as it was shown in Cheriet et al. (2007)). The results from the experiments reported here do not allow making direct comparisons with other calibration methods that use large calibration apparatus because of different evaluation measures. However, both Dumas et al. (2003) and Cheriet et al. (2007) had comparable performances to Aubin et al. (1997), which let us generalise the conclusions of the previous comparison. The major exception is that the method proposed in Cheriet et al. (2007) is able of handling eventual motion of the subject.

Comparison with the method proposed in Kadoury et al. (2010) that is based on estimating the geometrical parameters using a small calibration object (the same as in Kadoury et al. (2007b)) and then optimising them using high-level primitives automatically extracted from the radiographs, shows that the method proposed here achieves more accurate reconstructions in *in vitro* experiments (table 4.4). Differences are particularly noticeable at the pedicles landmarks. Additionally, the authors pointed out higher errors on the Z direction due to uncertainty in the depth axis, which is directly tackled by the method proposed here using the distance measuring device.

The errors presented here may increase on *in vivo* clinical conditions if subject positioning is not properly done. Additionally, for the complete version of the method, landmarks identification errors may be higher in areas where vertebrae are not so visible due to overlapping bone structures (e.g. upper thoracic ver-

tebrae on lateral radiographs). However, an *in vivo* validation was not possible because conventional 3D imaging techniques (e.g. CT, MRI) are unsuitable since they alter the spine configuration, and gold-standard calibration apparatus were not available.

Summarising, the method proposed here may be used in two particularly advantageous configurations. In the first, the rangefinder is used to help estimating the geometrical parameters, which are then directly used for performing the 3D reconstruction. In this version of the method, special care should be given to subject positioning, but there is no need for calibration objects. Additionally, no landmarks are necessary. The advantages are twofold: a source of error is eliminated, and user intervention is drastically reduced. One may argue that landmarks are still necessary for the reconstruction process; however, other reconstruction methods that are urging and that require much less supervision (e.g. Benameur et al. (2003, 2005); Dumas et al. (2008); Humbert et al. (2009); Kadoury et al. (2009a); Moura et al. (2009, 2010b)) may be used for accomplishing this task. As for the second version of the proposed method, the geometrical parameters are optimised and scale is corrected using a small calibration object. This version turned out to be more robust and accurate. However, it needs a considerably large set of points to be identified and subjects must wear a calibration object. This set of points is no larger than sets required by other methods Kadoury et al. (2007a,b) and there are semi-supervised methods that do not rely on calibration parameters that may be used for estimating these landmarks (e.g. Vaiton et al. (2004)). Additionally, to the best of our knowledge, the calibration object has the simpler geometry ever with the lowest impact on the contents of radiographs.

4.6 Summary

This chapter presented a novel method for the calibration of spine biplanar radiography that makes use of a distance that can be easily measured on site for providing 3D reconstructions with correct scale. This method is an adaptation to spine radiography of the method presented in the previous chapter, and addresses the definition of the acquisition protocol of radiographs among other problem-specific issues. The effects of the position, orientation and size of the calibration object were studied using numerical simulations with *in vivo* 3D data, and results enabled to state guidelines concerning these characteristics for best

results. Additionally, the effect of the precision of the distance measuring device was studied to determine its impact on reconstructions of the spine. Different variations of the method were analysed to identify the tradeoffs of skipping some of the calibration steps and to determine the gain of using a distance measuring device in the estimation of the calibration parameters. All these aspects were tested with pessimistic noise ranges to simulate conditions that may resemble conventional radiological setups where no rotatory platform is available. Results shown that the proposed method enables to improve reconstructions and that it enables to recover scale even when no calibration object is used and without requiring optimisation of the calibration parameters. Still, optimising parameters and correcting scale with a simple calibration object enables to achieve the best results. An *in vitro* validation was performed with a dried spine and results were consistent with simulations. Comparison with related work shown that the method proposed here has comparable or superior performance while introducing less artifacts in the contents of radiographs and being more suitable for standard clinical environments.

CHAPTER 5

Statistical-based 3D reconstruction of the spine

This chapter proposes a novel method for fast 3D reconstructions of the scoliotic spine from two planar radiographs (Moura et al., 2009, 2010b). The spine is represented in an articulated fashion for properly modeling inter-vertebral dependencies as well as their influence on vertebrae shape. We propose reconstructing the spine by deforming an articulated model towards fitting input data. An efficient use of the input enabled to achieve both fast and reliable reconstructions, making the method suitable for routine clinic.

This chapter starts with a description of statistical models of the spine. Then, the proposed method is described, as well as the validation protocol, which includes the evaluation of the accuracy of reconstructions as well as the accuracy calculating clinical indices by both non-experts and expert users. Finally, results are presented, discussed and compared with related work.

5.1 Statistical models of the spine

Statistical models of anatomical structures are often composed by a set of landmarks describing their geometry. Cootes et al. (1995) proposed aligning

point-based representations of several instances for computing their mean shape and principal deformations. These statistical models are known as Point Distribution Models (PDM), and were used intensively for describing shape variation of anatomical structures, such as the hippocampus (Styner et al., 2003), fiber tracts (Corouge et al., 2004), femoral heads (Rajamani et al., 2004), hands and faces (Vasconcelos and Tavares, 2008). For the case of the spine, a PDM could also be built from a set of 3D reconstructions obtained from supervised methods. However, the spine is a flexible structure where the position and orientation of vertebrae are not independent. Capturing the spine as a cloud of points does not differentiate vertebrae and, consequently, information is lost that may be important to capture spinal shape variability and vertebrae inter-dependencies.

For conveniently describing spine shape variability, Boisvert et al. (2008b) proposed using Articulated Models (AM). These models capture inter-vertebral variability of the spine geometry by representing vertebrae position and orientation as rigid geometric transformations from one vertebra to the other along the spine. Articulated models already have proven to be advantageous when only partial data about the shape of the spine is available. In Boisvert et al. (2008a), AM were used to infer 3D landmarks of vertebrae for which user input was missing. In this work, it was shown that it was possible to reconstruct the 102 three-dimensional landmarks that described the spine using only a quarter of the full set of manually identified point-matches. Additionally, in Boisvert et al. (2009), AM enabled inferring 3D reconstructions from a single radiograph, manually labelled with 102 landmarks. Still, to the best of our knowledge, AM were never used with very limited amount of input for reconstructing the spine. Additionally, reconstructions provided by the previous methods were achieved by searching for an optimal solution on the feature space of the AM, which may have 408 dimensions for a typical representation of the spine composed of 17 vertebrae with 6 points per vertebrae.

Here we propose using AM for reconstructing the spine from limited input while taking advantage of their representation, which models and restricts inter-vertebrae dependencies. For cooping with the large dimensionality of AM, we propose capturing the main deformation modes using Principal Components Analysis (PCA) (Jolliffe, 2002) and, thus, limiting the search space to the modes that explain most of the variation of the shape of the spine. In the next subsections, AM are described in more detail as well as the computation of the principal deformations from a sample of 3D reconstructions of the spine.

5.1.1 Articulated representations of the spine

In an articulated representation of the spine, only the first vertebra (i.e. L5) has an absolute position and orientation, and the following vertebrae are dependent from their predecessors (figure 5.1):

$$T_i^{abs} = T_1 \circ T_2 \circ \cdots \circ T_i, \quad \text{for } i = 1..N, \tag{5.1}$$

where T_i^{abs} is the absolute geometric transformation for vertebra i, T_i is the geometric transformation for vertebra i relative to vertebra $i-1$ (with the exception of the first vertebra, which is related to the world reference frame), $\circ$ is the composition operator, and N is the number of vertebrae represented by the model.

Source: adapted from Boisvert et al. (2008c)

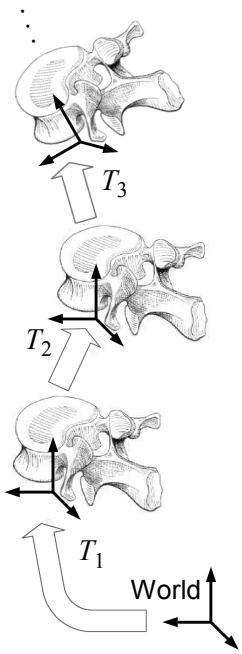


FIGURE 5.1: The spine as an articulated model: the position and orientation of each vertebra is expressed by inter-vertebrae geometric transformations.

In order to include data concerning vertebrae morphology, a set of landmarks is expressed in the local coordinate system of each vertebra. The absolute coordinates for each landmark may be calculated using the equation:

$$p_{i,j}^{abs} = T_i^{abs} \star p_{i,j}, \quad \text{for } i = 1..N, j = 1..M, \tag{5.2}$$

where $p_{i,j}^{abs}$ are the absolute coordinates for landmark j of vertebrae i, $p_{i,j}$ are the relative coordinates, $\star$ is the operator that applies a transformation to a point, and

M is the number of landmarks per vertebra. Thus, an articulated representation of the spine that models both global and local shape may be expressed as a vector that includes inter-vertebral rigid transformations and relative landmarks for each vertebra:

$$s = [T_1, p_{1,1}, \ldots, p_{1,M}, \ldots, T_N, p_{N,1}, \ldots, p_{N,M}]. \tag{5.3}$$

5.1.2 Statistics on articulated models of the spine

For building a statistical model of the spine, a set of 3D reconstructions are first represented on an articulated fashion (see equation 5.3), making a set of articulated representations S. Calculating statistics for such a set is not trivial because AM include geometric transformations, which do not belong to an Euclidian space. Therefore, statistics like the arithmetic mean cannot be directly computed. However, rigid transformations belong to a Riemannian manifold (Pennec, 2006). Hence, Riemannian manifolds can be used to map articulated representations on Euclidian space, which allow the computation of centrality and dispersion statistics, namely the Fréchet mean (μ) and the generalized covariance matrix (Σ) (Boisvert et al., 2008b; Pennec, 2006).

Let $\mathcal{M}_i$ be the manifold that maps articulated representations for a given vertebral level i on Euclidian space. Then, the Fréchet mean for that vertebral level i is the element μ_i of the manifold $\mathcal{M}_i$ that minimises the sum of distances to all elements x_m of the same manifold:

$$\mu_i = \underset{x \in \mathcal{M}_i}{\arg\min} \sum_m d(x, x_m)^2, \quad \text{for } i = 1..N, \tag{5.4}$$

where d is the Riemannian metric that defines the distance between two vertebrae's representations. Let $a_i = \{T_a, p_a\}$ be the representation of vertebral level i of spine a defined by a relative geometric transform T_a and a set of relative landmarks p_a, and $b_i = \{T_b, p_b\}$ a representation of vertebral level i of spine b. The distance between the two vertebrae is given by:

$$d(a_i, b_i) = d(\{T_a, p_a\}, \{T_b, p_b\}) = D_\beta(T_b^{-1} \circ T_a, p_b - p_a), \tag{5.5}$$

with

$$D_\beta(\Delta T, \Delta p)^2 = D_\beta(\{\omega, t\}, \Delta p)^2 = \|\beta\omega\|^2 + \|t\|^2 + \|\Delta p\|^2, \tag{5.6}$$

where $t \in \mathbb{R}^3$ and $\omega \in \mathbb{R}^3$ are respectively the translation and the rotation in axis/angle form that make the geometric transform ΔT (extracted in the same

way as described in section 3.2 with Rodrigues' formula), and $\beta \in \mathbb{R}$ is a weight that balances the impact of rotation in the metric with respect to translation and landmarks' coordinates. Having the mean calculated for all vertebral levels, $\mu = [\mu_1, \ldots, \mu_N]$, it is possible to calculate the departures from the mean articulated model μ to each of the articulated representations of S. Let ΔS be the departures from the mean, the covariance matrix may be calculated as:

$$\Sigma = \frac{\Delta S^T \Delta S}{Q}, \tag{5.7}$$

where Q is the number of spines represented in S. The reader is addressed to Pennec (2006) for more details on statistics on Riemannian manifolds and to Boisvert et al. (2008b) for their application to articulated models of the spine.

5.1.3 Deformable articulated model of the spine

The method proposed here uses an articulated model of the spine comprised of $N = 17$ vertebrae (from L5 to T1) and $M = 6$ landmarks per vertebra. The first two landmarks are the centre of the superior and inferior endplates ($j = 1..2$) and the remaining four are the superior and inferior extremities of both pedicles ($j = 3..6$). The origin of each vertebra coordinate system is located at the centre

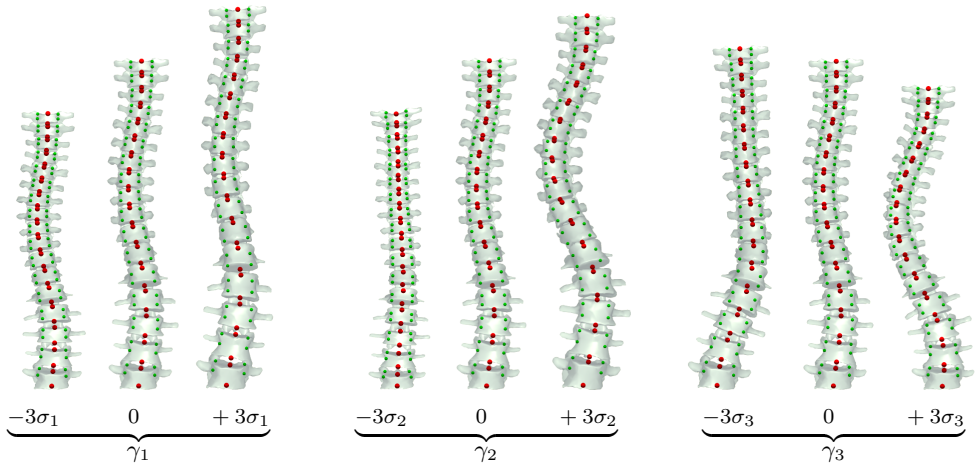


FIGURE 5.2: Effect of varying the weight (γ_i) of each of the first three principal deformation modes of the statistical shape model in turn for -3, 0 (mean model), and 3 times the standard deviation (σ_i) of the deformation mode. The statistical shape model describes the variability of 6 landmarks per vertebrae (endplates – red strong points medially located; pedicles – green small points laterally located) by modeling their relative location, orientation and shape on an articulated fashion. For illustration purposes, 3D models of complete vertebrae were rendered.

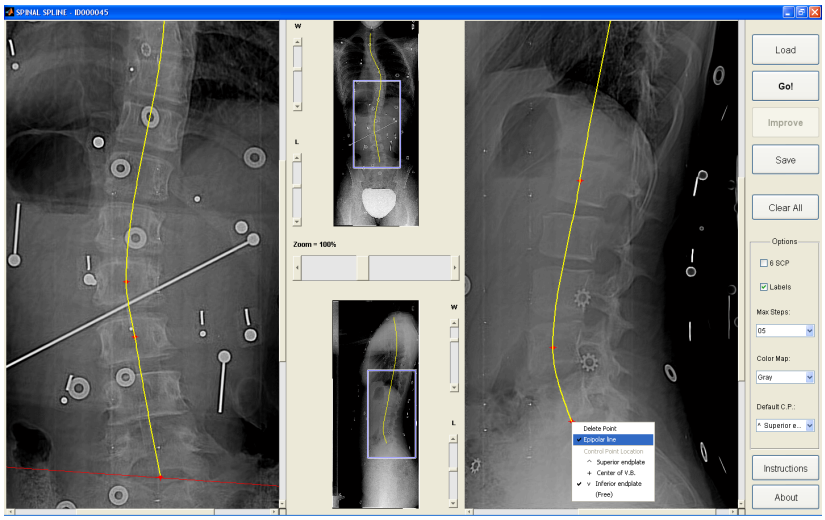


Figure 5.3: Graphical user interface (GUI) designed for identifying the splines: this figure illustrates how the GUI may provide epipolar lines; in this case, an epipolar line is drawn on the frontal radiograph (red line) that corresponds to the control point identified on L5 vertebra on the lateral radiograph.

of the pedicles' landmarks. This representation was used due to the availability of reconstructions with theses features, and because it remains one of the standard 3D representations of the spine.

For reducing the number of dimensions of the articulated model while capturing the main deformation modes, a Principal Components Analysis (PCA) (Jolliffe, 2002) was performed on the covariance matrix Σ (figure 5.2). Using PCA in a linearised space, an articulated representation of a spine s may be generated by linearly combining the eigenvectors of the covariance matrix and then by composing the result with the mean articulated model:

$$s = \mu \circ \sum_i \gamma_i v_i \tag{5.8}$$

where μ is the Fréchet mean of a sample of articulated spine models, γ_i is the weight associated with the i^{th} eigenvector of the covariance matrix, and v_i is the i^{th} eigenvector of the covariance matrix. The composition between the mean model and the variation of the principal components may be formally described as:

$$\begin{aligned} s = \mu \circ \Delta = \; [\; & T_1 \circ \Delta T_1 \; , \; p_{1,1} + \Delta p_{1,1} \; , \; \ldots \; , \; p_{1,M} + \Delta p_{1,M} \; , \; \ldots \; , \\ & T_N \circ \Delta T_N \; , \; p_{N,1} + \Delta p_{N,1} \; , \; \ldots \; , \; p_{N,M} + \Delta p_{N,M} \;]. \end{aligned} \tag{5.9}$$

Thus, PCA allows the generation of articulated representations of the spine by manipulating the principal deformation modes, that is, the values of γ_i. Moreover, limiting the range of values γ_i allows to ensure valid reconstructions (see section 5.3). After generating an articulated representation s, a 3D reconstruction may be obtained by first calculating the absolute transformation of each vertebrae (equation 5.1) and then the absolute 3D position of every landmark (equation 5.2).

5.2 Spline guided deformation of the articulated model

The method proposed here deforms a statistical model of the spine (section 5.1.3) towards fitting the spine midline captured by two calibrated radiographs (typically one frontal and one lateral). User input is limited to placing a few control points for identifying the spine midline in the two radiographs using parametric splines. Both splines should begin at the centre of the superior endplate of vertebra T1 and should end at the centre of the inferior endplate of L5. These are the only stereo-correspondent points that are required. A Graphical User Interface (GUI) was developed that enables users to identify the splines and helps ensuring stereo-correspondence of the endpoints of the splines by displaying epipolar lines (figure 5.3). For enabling full use of the input data, users may place the remaining control points at specific anatomical positions, i.e. centre of superior endplates, centre of inferior endplates, or centre of vertebral bodies. Placing control points at particular anatomical features provides input concerning vertebrae location that is used for improving reconstruction accuracy (section 5.4). Typically, for faster interaction, users place all control points at a default location, i.e. the centre of superior endplates.

Three-dimensional reconstructions of the spine based on the user-defined splines are achieved using an optimisation process that iteratively deforms the statistical model towards minimising the distance between the projected landmarks of the reconstructed model and the splines. Thus, the goal of the optimisation process is to find the values of γ_i (equation 5.8) that generate the spine configuration s that best fits the user-defined splines of both radiographs. The fitting error was defined as the distance between (a) the absolute position of the endplates of the deformed model and (b) the user-defined splines. For calculating such distance, the endplates of the deformed model are first projected to both radiographs (e.g. frontal and lateral). Then, for each radiograph, the coordinates

of the projected endplates (p^{2D}) have to be mapped to the user-defined spline in order to calculate the distance between endplates and the spline (figure 5.4). The mapped locations $u = \{x_u, y_u\}$ are calculated in the following way:

- y_u are obtained using linear interpolation: the values of the y coordinate of the projected endplates are scaled to fit the height of the user-defined spline;
- x_u are the values of the x coordinate on the user-defined spline at y_u, which are found using piecewise cubic interpolation (Fritsch and Carlson, 1980), assuming that the y coordinate along the spine midline is monotonically increasing.

An alternative capable of handling non-monotonicity was proposed in Moura et al. (2009), which is based on using one-dimension coordinates along the spline rather than the y coordinate directly. However, the assumption of monotonicity is generally valid for the case of the spine, and the solution described here shows comparable results while being more computationally efficient than the one described in Moura et al. (2009).

The cost function may now be defined as:

$$C = \sum_{i=1}^{N} \sum_{j=1}^{2} \sum_{k=1}^{2} \left\| p_{i,j,k}^{2D} - u_{i,j,k} \right\|^2, \tag{5.10}$$

where $p_{i,j,k}^{2D}$ is the projection of the 3D endplates ($p_{i,j}^{abs}$) to radiograph k, and $u_{i,j,k}$ are their estimated locations on the user-defined spline of the same radiograph. Minimising function C is a nonlinear least-squares problem, which is solved with a trust-region-reflective algorithm (Coleman and Li, 1996). This algorithm, previously described in section 3.5.2, requires an initial solution that is iteratively updated until convergence. The initial solution was chosen to be $\gamma_i = 0$ for all principal components, i.e. the mean model of the spine.

5.3 Generating plausible spine configurations

The weights γ_i are limited to an hyperellipsoid in the parameter space such that $|\gamma_i| \leq 3\sqrt{\lambda_i}$, being λ_i the eigenvalues of the covariance matrix (Σ) (Cootes et al., 1995). In other words, departures from the mean shape are limited to three standard deviations to avoid outliers. Moreover, the cost function was modified to include a term that promotes models that are more likely with respect to the prior model. This is done using the Mahalanobis distance (Maesschalck et al.,

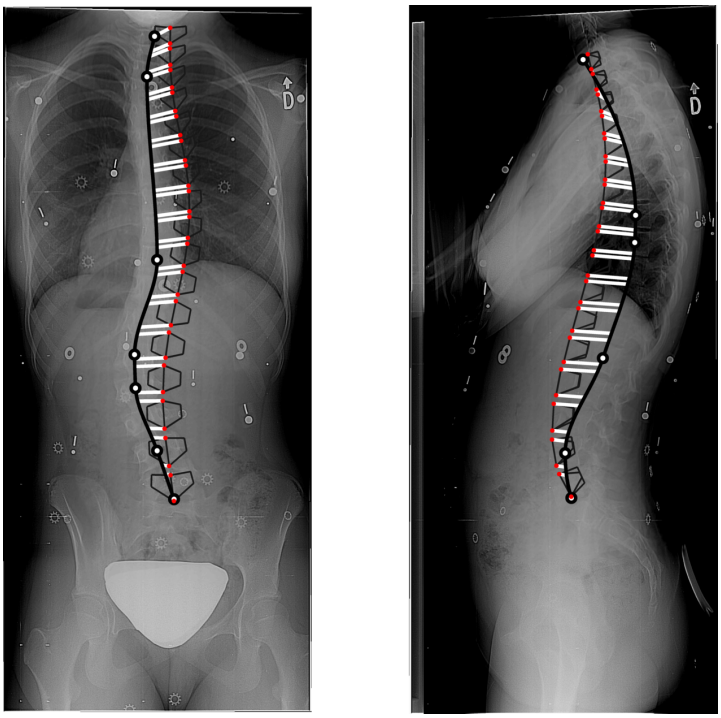


FIGURE 5.4: Fitting error of the deformable AM, calculated as the distance between the endplates of the AM and their estimated positions on the user-defined splines. The AM is first projected to the PA and LAT radiographs where the operator identified the splines and then the error (white thick line-segments) is calculated for each endplate on each radiograph. The AM is represented by 6 points per vertebra connected using black thin line-segments, and the user-defined splines are represented by thick black curves with control points as white circles.

2000) on the feature space of the articulated model, which is defined as:

$$D = \sqrt{\Delta^T \Sigma^{-1} \Delta}, \tag{5.11}$$

where Δ represents the deformation, i.e. the variation to the mean like defined in equations 5.8 and 5.9. Then, the cost function becomes:

$$C = \sum_{i=1}^{N} \sum_{j=1}^{2} \sum_{k=1}^{2} \left\| p_{i,j,k}^{2D} - u_{i,j,k} \right\|^2 + (\alpha D)^2 , \tag{5.12}$$

where α is used for balancing the weight of the prior spine shape knowledge with respect to the spline fitting error. The value of parameter α essentially depends of the pixel size of the radiographs since the spline error is computed in pixels.

5.4 Refinement of vertebrae location

The fitting process just described captures the shape of the spine by placing vertebrae on their probable location along the spine midline, which may not be the correct one since there might be a range of valid solutions. For improving spine reconstructions without requesting additional information to the user, the location of the control points of the splines are used. However, despite control points being placed at specific anatomical positions (like described in section 5.2), it is not known on which vertebrae they lie. For tackling this issue, the two nearest vertebrae of the AM are selected as candidates for each control point after a first minimisation of equation 5.12. Then, the nearest candidate is elected if the level of ambiguity is low enough. This may be formalised on the following way:

$$\frac{d_{m,1}}{d_{m,2}} \leq \varepsilon, \quad \text{with } d_{m,1} \leq d_{m,2} \tag{5.13}$$

where $d_{m,1}$ is the distance of control point m to the nearest candidate of the AM, $d_{m,2}$ is the distance to the second nearest candidate, and ε is a threshold that defines the maximum level of ambiguity allowed. Since $d_{m,1} \leq d_{m,2}$, ambiguity has maximum value of 1 (one) when the candidates are equidistant to the control point, and minimum value of 0 (zero) when the nearest candidate is in the exact location of the control point. The list of candidates for a given control point depends of the anatomical position where it was placed, e.g. if a control point was placed on the superior endplate, only the superior endplates of the AM would be candidates for that point.

After determining the set of elected candidates, the optimisation process is repeated, but now including a third component that is added to equation 5.12. This term attracts the elected vertebrae of the articulated model towards their corresponding control points (figure 5.5). Let E be the set of elected candidates, the cost function may be redefined as:

$$C = \sum_{i=1}^{N} \sum_{j=1}^{2} \sum_{k=1}^{2} \left\| p_{i,j,k}^{2D} - u_{i,j,k} \right\|^2 + (\alpha D)^2 + \sum_{m \in E} \|d_{m,1}\|^2 . \tag{5.14}$$

When the second optimisation finishes, the vertebrae location of the articulated model should be closer to their real position, and some of the ambiguities may be solved. Therefore, several optimisation processes are executed iteratively while the number of elected candidates increases.

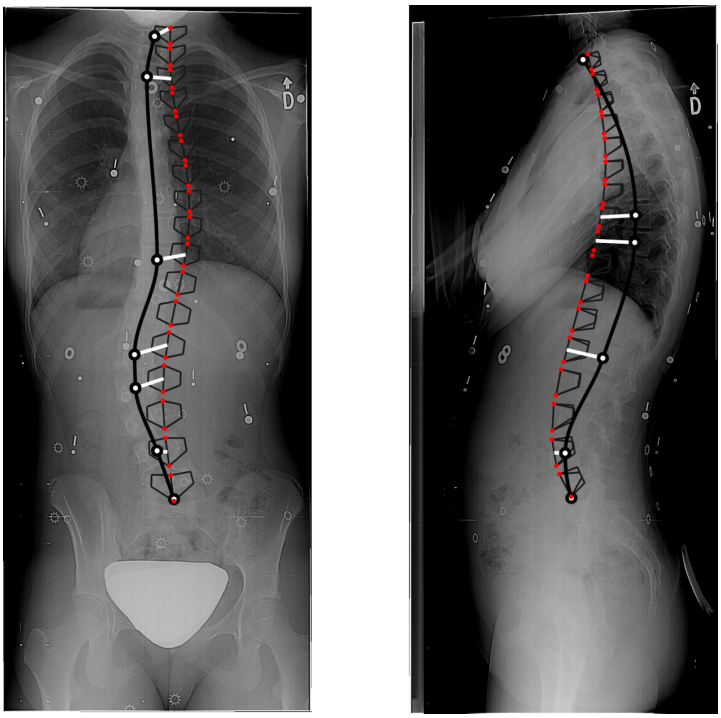


FIGURE 5.5: Using the location of the control points of the splines to improve fitting: since users place each control point over a vertebra, the location of the control point on the radiograph is used to attract the nearest vertebra of the AM. In this illustration all control points were placed at the centre of the vertebral bodies (with the exception of T1 and L5). The AM is represented by 6 points per vertebra connected using black thin line-segments, the user-defined splines are represented by thick black curves with control points as white circles, and the distance between control points and their nearest vertebra on the AM is represented as white thick line-segments.

Concerning the value of ε, using a low threshold of ambiguity may result in a considerable waste of control points due to an over-restrictive strategy. On the other hand, a high threshold of ambiguity may produce worst results, especially when there are control points placed at erroneous locations. For overcoming this issue, a dynamic thresholding technique is used that begins with a restrictive threshold where only candidates that are at half the distance to the target or less than the second nearest candidate are elected ($\varepsilon = 0.50$). Then, when no candidates are elected, ambiguity is relaxed (by increments of 0.10) up to a maximum threshold ($\varepsilon = 0.70$). If any control points remain ambiguous at this stage, they are considered to be unreliable.

5.5 Method evaluation

The validation study uses radiographs from scoliotic patients that were acquired at Saint-Justine Hospital Centre in Montreal, Canada, with a FCR7501 system (Fuji Medical, Tokyo, Japan), producing 12-bits grayscale digital images with resolution of 2140×880 pixels. Two radiographs were available for each examination, one Posterior-Anterior (PA) and one Lateral (LAT). Patients positioning and radiography calibration was ensured by the system proposed in Cheriet et al. (2007).

Parameter $\alpha = 2.5$ was empirically found to provide a good balance between the weight of the prior spine shape knowledge with respect to the other terms of the cost function. Since α depends of the pixel size of the radiographs this value may not be adequate for other systems.

Accuracy was evaluated by comparison with reconstructions from a previously validated method (Delorme et al., 2003), which is based on the manual identification of 6 stereo-corresponding points by an experimented operator followed by triangulation for finding their 3D positions.

The deformable articulated model was calculated with $\beta = 1.0$ using 295 3D spine reconstructions (Cobb angle in the interval $[4°, 86°]$) that did not included any of the patients of the testing sets. For enhancing computational performance, a different number of principal components were used depending on the stage of the reconstruction method. In the last optimisation (e.g. when there are no ambiguous control points), the principal components that explain 99% of the spine shape variation were used, and in the previous optimisations only 95% of the components were used.

5.5.1 Experiments with experimented users

The first set of experiments measured accuracy of spine reconstruction, vertebrae location and rotation, and selected clinical indices for a total of 30 patients: 10 moderate idiopathic scoliosis with Cobb angle in the interval $[22°, 43°]$ and mean value of 33°, and 20 severe idiopathic scoliosis with Cobb angle in the interval $[44°, 70°]$ and mean value of 55°. Reconstructions were performed with the method proposed here by an experimented operator. Additional, the relation between the amount of user interaction and the accuracy of the reconstruction was studied with simulated splines.

User interaction versus reconstruction accuracy

For determining the influence of the amount of user input on the reconstruction accuracy, splines were automatically generated with increasing number of control points from 5 to 17 per radiograph. In particular, for each patient and for each view, the manually identified endplates of the reference data were used to build the input spline by selecting as control points the superior endplates that minimised the 2D distance of every reference endplate to the spline.

Spine reconstruction accuracy

Spine reconstruction accuracy was evaluated as the point-to-point euclidian distance between each reconstructed point and its location on the reference data. RMS 3D reconstruction errors were first calculated for each exam and for both endplates and pedicles. The mean, standard deviation and maximum RMS errors for each test set were calculated. Results for patients with moderate scoliosis were compared with the values presented in Kadoury et al. (2009a) where the same statistics were used for a similar sample of patients.

Influence of each term in the cost function

In order to determine the influence of the terms of the cost function (equation 5.14), the 3D reconstruction errors of the previous experiment using all terms were compared with results obtained by not including either the Mahalanobis distance, the vertebrae location refinement, or both.

Vertebrae location and orientation accuracy

Accuracy of vertebrae location and orientation was measured as the Root Mean Square of the Standard Deviation (RMS_{SD}) of the error between reconstructions with the proposed method (observation 1) and the reference data (observation 2), as proposed in Dumas et al. (2008):

$$RMS_{SD} = \sqrt{\frac{\sum_m \left(\frac{\sum_n (\bar{\alpha} - \alpha_n)}{n}\right)^2}{m}}, \tag{5.15}$$

where $\bar{\alpha}$ is the mean of the $n = 2$ observations, and m is the number of computed locations or orientations about either of the axes. A vertebral reference frame

was associated to each vertebra based on the definition of Stokes and the Scoliosis Research Society (Stokes, 1994) and was used to assess location and orientation: origin of the reference frame at the centre of the vertebral body and calculated as the midpoint between the centres of the superior and inferior endplates; Z axis passing by the centre of both endplates; Y axis parallel to a line that passes by the centre of the left and right pedicles; X axis perpendicular to the remaining axes. Results for the moderate scoliosis testing set were compared with Dumas et al. (2008).

Clinical indices accuracy

Accuracy of the proposed method measuring clinical indices was evaluated as the mean and standard deviation of the differences to the reference method for the following indices: Cobb angle on the PA, Cobb angle on the plane of maximum deformity, orientation of the plane of maximum deformity, kyphosis and lordosis. Additionally, a Wilcoxon signed-rank test was performed for identifying if there were significant differences ($p \leq 0.05$) between results of the two methods.

Reconstruction time

Finally, reconstruction time was evaluated by measuring the user interaction time needed for identifying the splines as well as the computation time for delivering the 3D reconstruction. Average times were computed for the two testing sets and comparisons were made with other spline-based methods (Dumas et al., 2008; Humbert et al., 2009; Kadoury et al., 2009a). Computation times were measured on a Desktop PC with an AMD Phenom II X2 550 3.10 GHz processor and 2 GBytes of memory for an implementation in Matlab of the proposed method.

5.5.2 Experiments with non-expert users

Accuracy of the 3D reconstruction was also evaluated for non-expert users on a total of 14 patients with idiopathic scoliosis (Cobb angle in the interval $[29°, 53°]$ and mean value of $40°$). Reconstructions were performed with the method proposed here by two volunteers with limited knowledge of spine radiology. Both had 20 minutes of training with the software tool before performing the experiment. Input errors were calculated by computing the distance of the control

points to both the location of the control points on the reference data, and to the nearest point of the spine. Accuracy was evaluated as the point-to-point euclidian distance between each reconstructed point and its location on the reference data. Additionally, the same clinical indices of the previous experience were evaluated for computing both accuracy and inter-observer variability. The Wilcoxon signed-rank test was performed for identifying if there were significant differences ($p \leq 0.05$) to the reference method, as well as differences between the results of the two operators. Interaction time was also recorded.

5.6 Results

This section starts with results of the method for splines identified by an experimented operator, and the last subsection is dedicated to results by non-expert operators.

5.6.1 Influence of each term in the cost function

Results show that adding the Mahalanobis distance term to the spline fitting term enabled to decrease 3D reconstruction errors by 10% on the endplates and 8% on the pedicles while adding the distance to the control points decreased errors by 55% on the endplates and 35% on the pedicles (figure 5.6). Both terms were always able to decrease reconstruction errors but the impact of using input concerning control points was always superior, particularly in the endplates.

5.6.2 User interaction versus reconstruction accuracy

Results show that the reconstruction accuracy of the endplates increases with the number of control points of the splines, on both moderate and severe scoliosis (figure 5.7). This is particularly observable until a given limit where saturation occurs (e.g. 11 control points on moderate scoliosis). Reconstructions with 5 control points did not conveniently described the spine midline of 3 of the 20 patients with severe scoliosis. This resulted on 3 bad reconstructions with ambiguities mapping control points to the articulated model, which produced higher reconstruction errors. In all other tests all ambiguities were completely and correctly solved. As for the pedicles, no considerable improvement is observed by

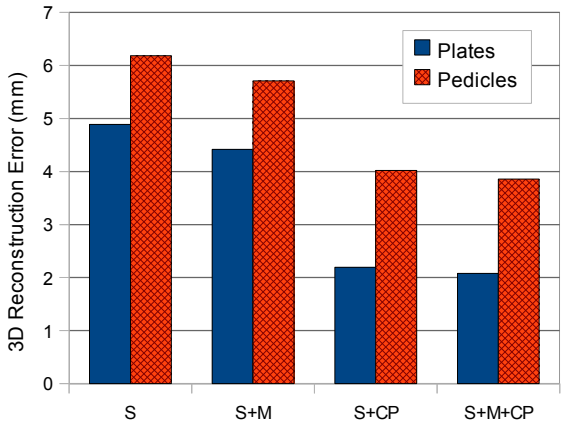


Figure 5.6: Impact on the 3D reconstruction accuracy of the terms of the cost function; S – distance to splines (equation 5.10); M – Mahalanobis distance (second term of equation 5.12); CP – distance to control points (third term of equation 5.14).

increasing the number of control points beyond the minimum necessary to describe the spine midline (≥ 7 for severe scoliosis). Comparison between the two test sets, for a number of control points ≥ 7, revealed that reconstruction errors on the endplates were on average 0.2 mm superior on the severe test set while on the pedicles this difference was of 0.5 mm.

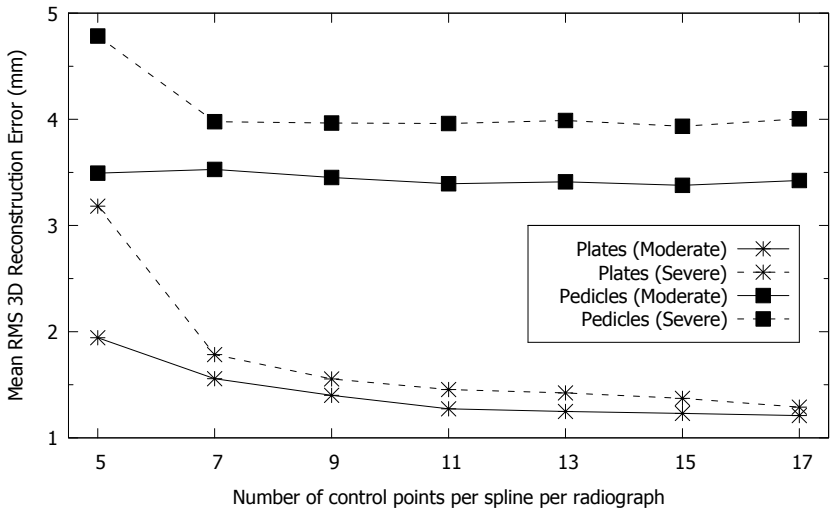


Figure 5.7: Mean RMS 3D reconstruction error vs number of control points per spline and per radiograph.

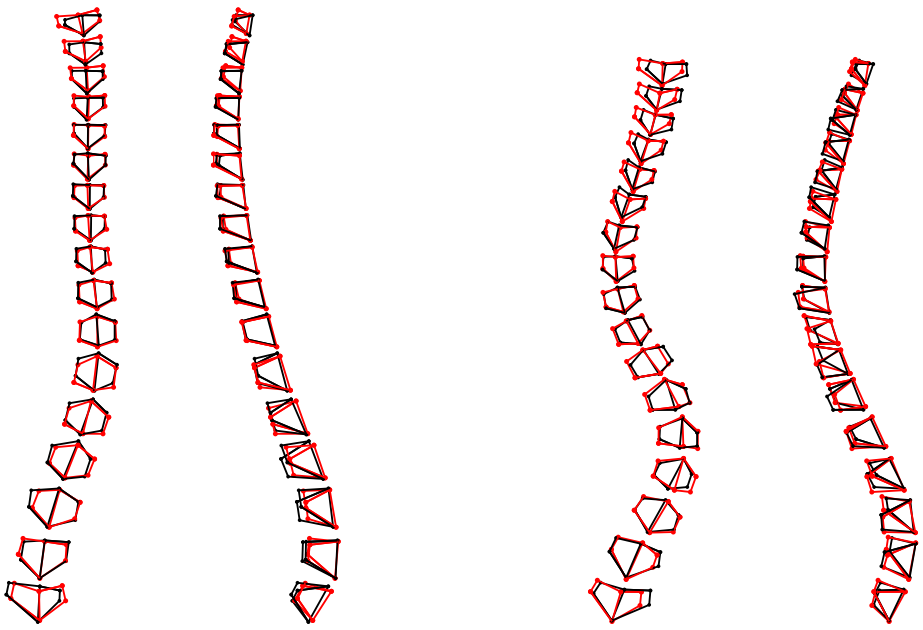


FIGURE 5.8: Comparison of 3D reconstructions using the proposed method by an experimented operator (at red) vs the reference data (at black). On the left, PA and LAT views of the median case of the mild scoliosis group and, on the right, PA and LAT views of the median case of the severe scoliosis group.

5.6.3 Spine reconstruction accuracy

Results with an experimented operator show average RMS reconstruction errors of 2.0 mm and 2.1 mm on the endplates' landmarks for moderate and severe scoliosis respectively (table 5.1). Reconstruction errors of the pedicles were superior on severe scoliosis (4.0 mm) when compared with moderate scoliosis (3.5 mm). Pedicles' reconstruction errors were superior to endplates' reconstruction errors on all patients. The maximum reconstruction error was observed on the patient with the highest Cobb angle.

The operator identified an average of 7 control points on both PA and lateral radiographs when reconstruction patients with moderate scoliosis, and an average of 9 control points on the PA and 7 on the lateral for severe scoliosis. The majority of control points were placed on the centre of the superior endplate and only on a few vertebrae the operator choose to place control points on the bottom endplate. The proposed method was always able to solve the ambiguities when mapping control points to vertebrae of the articulate model, and therefore all control points were always used for refining vertebrae location. Figure 5.8 illustrates the differences between 3D reconstructions of the proposed and refer-

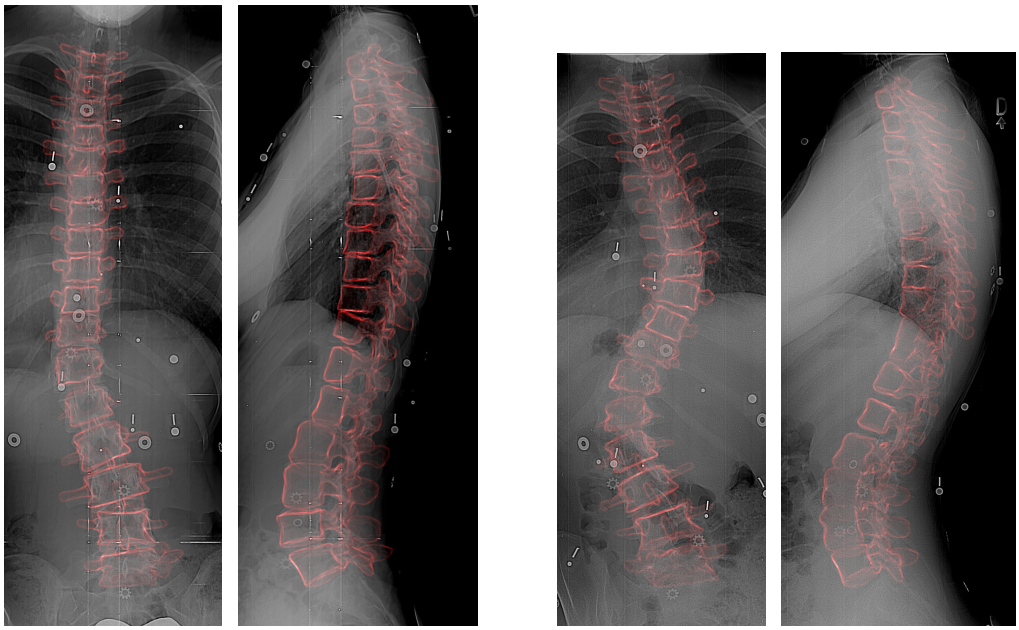


Figure 5.9: Projection on the radiographs of 3D vertebrae models deformed using the 6 points per vertebrae determined by the method proposed here. On the left, PA and LAT views of the median case of the mild scoliosis group and, on the right, PA and LAT views of the median case of the severe scoliosis group.

ence methods for the median case of each group of patients, and figure 5.9 shows the projection on the radiographs of 3D vertebrae models deformed using the 6 points per vertebrae that are calculated by the proposed method (for details concerning the deformation algorithm please consult appendix A).

5.6.4 Vertebrae position and orientation accuracy

Results for this experiment are presented on table 5.2 and show that the proposed method presents comparable accuracy locating vertebrae on patients with moderate and severe scoliosis. Regarding orientation, the same was observed for rotations about all axes with the exception of axial rotation (Z axis) where the error on severe scoliosis (4.4°) was superior to moderate scoliosis (3.3°).

5.6.5 Clinical indices accuracy

Results for this experiment are presented on table 5.3 and show that there are no evidences of significant differences (with $p \leq 0.05$) for all evaluated clinical indices on patients with moderate scoliosis. On severe scoliosis, Cobb angle at

the maximum plane of deformation and Kyphosis shown statistically significant difference, both with $p = 0.03$.

5.6.6 Reconstruction time

Average reconstruction time, for both computation and interaction by an experimented user, are presented in table 5.4 and show that computation time is the same for the two sets, while the user required on average more 20 seconds for identifying the spine midlines on patients with severe scoliosis.

5.6.7 Results for non-expert users

Errors on the 3D reconstruction of the anatomical landmarks by non-expert users are superior to the previously presented results by an experimented operator (table 5.5). Interaction times were on average 1.5 minutes, with user A being 10 seconds faster on average than user B. User A used on average 6 control points on each radiograph and the input errors identifying the spline and the control points were comparable. User B utilised on average 5 control points per spline and errors identifying the exact place of the control points were superior to errors describing the spine midline. The method was able to solve all ambiguities affecting control points for user B, while for user A in 4 of the 14 exams the method was not able to affect one of the control points to a vertebra.

Evaluation of the clinical indices shows that despite of the erroneous input and inferior 3D reconstruction accuracy, the method is still able of providing the calculation of clinical indices with no evidence of significant differences (for $p \leq 0.05$) to the reference method (table 5.6). The exception was Lordosis for user A with $p = 0.03$. The inter-observer variability study shows that only for Lordosis the results of the clinical indices were also significantly different ($p = 0.02$).

5.7 Discussion

Results shown that the method proposed here is able to provide accurate clinical indices with both expert and non-expert users in less than 2 minutes. To the best of our knowledge, this is the first study that shows that fast and useful reconstructions of the spine may be obtained on-the-fly by non-experts. In addition, expert knowledge increases reconstructions accuracy outperforming

Table 5.1: RMS reconstruction errors obtained by an experimented operator for the Severe and Moderate scoliosis test sets, and comparison with Kadoury et al. (2009a).

Method	N	Cobb angle [min−max]	Mean±SD [max] 3D error (mm)	
			Plates	Pedicles
Proposed	20	[44 − 70°] (severe)	2.1 ± 0.3 [2.9]	4.0 ± 0.9 [6.1]
Proposed	10	[22 − 43°] (moderate)	2.0 ± 0.3 [2.3]	3.5 ± 0.4 [4.3]
Kadoury (2009)	20	[15 − 40°] (moderate)	2.2 ± 0.9 [4.7]	2.0 ± 1.5 [5.5]

Table 5.2: RMS_{SD} location and orientation errors by an experimented operator for the Severe and Moderate scoliosis test sets, and comparison with Dumas et al. (2008) (orientation is expressed as a rotation about the given axis).

Method	N	Mean Cobb angle	Location (mm)			Orientation (°)		
			X	Y	Z	X	Y	Z
Proposed	20	55° (severe)	0.6	0.5	0.5	1.2	1.5	4.4
Proposed	10	33° (moderate)	0.5	0.5	0.4	1.2	1.3	3.3
Dumas et al. (2008)	11	30° (moderate)	0.9	1.1	2.3	1.3	2.0	3.2

Table 5.3: Average differences on clinical indices between the proposed method by an experimented operator and the reference method, and results of the Wilcoxon signed-rank test: NS – no evidence of significant differences, S – significant differences with $p \leq 0.05$ (Cobb_{PA} – Cobb angle on the PA, $\text{Plane}_{\text{Max}}$ – Plane of maximum deformity, Cobb_{Max} – Cobb angle on the $\text{Plane}_{\text{Max}}$).

	Moderate		Severe	
Index	Mean ± SD diff.	p	Mean ± SD diff.	p
Cobb_{PA}(°)	0.4 ± 3.1	0.85 (NS)	1.0 ± 2.4	0.09 (NS)
Cobb_{Max}(°)	0.7 ± 4.9	0.70 (NS)	1.6 ± 2.8	0.03 (S)
$\text{Plane}_{\text{Max}}$(°)	2.7 ± 17.7	0.64 (NS)	1.5 ± 16.2	0.57 (NS)
Kyphosis(°)	-1.2 ± 5.7	0.92 (NS)	0.9 ± 1.6	0.03 (S)
Lordosis(°)	2.2 ± 4.6	0.16 (NS)	0.0 ± 5.8	0.91 (NS)

Table 5.4: Comparison of the average reconstruction time (min:s) with other spline-based methods for patients with moderate scoliosis (statistics for severe scoliosis are included inside brackets when available).

	Proposed	Humbert et al. (2009)	Kadoury et al. (2009a)	Dumas et al. (2008)
User interaction	1:30 [1:50]	2:30 [3:00]	n.a.	5:00
Computation	0:03 [0:03]	0:04 [0:04]	2:24	n.a.

TABLE 5.5: Mean 3D reconstruction errors for splines identified by non-experts with 20 minutes of training. Mean values (± S.D.) for input and reconstruction errors. C.P. refers to the control points used for identifying the splines.

User	Input error (pixels)		Reconstruction Error (mm)	
	Spline	C.P.	Endplates	Pedicles
A	6.3±5.7	5.9±5.4	3.5±2.0	4.8±2.5
B	5.2±5.0	8.7±8.6	3.3±2.1	4.5±2.5

TABLE 5.6: Average differences on clinical indices between the proposed method by two non-expert operators and the reference method, and between the two operators. Results of the Wilcoxon signed-rank test are also presented and significant differences ($p \leq 0.05$) are marked with an asterisk. Cobb_{PA} – Cobb angle on the PA; $\text{Plane}_{\text{Max}}$ – Plane of maximum deformity; Cobb_{Max} – Cobb angle on the $\text{Plane}_{\text{Max}}$.

	User A vs Reference		User B vs Reference		User A vs User B	
Index	Mean ± SD	p	Mean ± SD	p	Mean ± SD	p
$\text{Cobb}_{\text{PA}}(^\circ)$	1.7 ± 3.8	0.11	-0.2 ± 2.7	0.86	-2.0 ± 5.0	0.14
$\text{Cobb}_{\text{Max}}(^\circ)$	2.5 ± 6.5	0.17	1.2 ± 4.5	0.33	-1.4 ± 5.2	0.27
$\text{Plane}_{\text{Max}}(^\circ)$	-2.4 ± 24.7	0.73	7.1 ± 24.3	0.16	9.4 ± 20.9	0.17
Kyphosis(°)	-0.8 ± 5.2	0.60	-0.2 ± 4.9	0.50	0.6 ± 2.3	0.27
Lordosis(°)	-5.9 ± 9.1	0.03*	-1.1 ± 7.3	0.63	4.8 ± 7.2	0.02*

previous methods locating vertebrae and determining their orientation, while requiring less reconstruction time. Despite computation times being among the fastest for 3D reconstructions of the spine, computation time may be considerably decreased by calculating the derivatives of the cost function analytically, instead of using finite differences. This would enable users to have reconstructions as they identify control points, which would allow to interactively refining reconstructions.

Results show that making a complete use of the user input, in particular, by utilising the location of the control points has an evident impact on the reconstruction quality. Additionally, identifying more control points on the splines enables to improve the accuracy of the reconstruction of the endplates' centres, creating a tradeoff between reconstruction accuracy and interaction time that may be adjusted according to users' needs and the objective of the examination. In fact, results with simulated splines show that reconstruction accuracy of the endplates could be superior if the operators had chosen to identify more control points, although this additional interaction may have no effect on the accuracy of the reconstruction of pedicles.

Using an optimisation scheme is a key feature of the method since it allows to easily incorporate variable number of control points per spline, that may be placed

on any vertebra. Several previous methods rely on regression or interpolation and either require having a fixed number of control points (e.g. Vaiton et al. (2004)), or additional data has to be placed at predefined vertebrae independently of its visibility while obliging to locate them (e.g. Humbert et al. (2009)), or simply ignore the position of the control points (e.g. Kadoury et al. (2009a)). Using an optimisation scheme also allows to level the weight of the prior with respect to the user input, which may be useful for different kinds of operators, e.g. non-experts may benefit from having more weight on the prior since the input is less reliable, while for experts the weight of the prior may be lower. Another interesting feature that was not explored in this study is that the method can easily be generalised to handle a different number of radiographs. Using more than two radiographs should not have considerable impact on the 3D reconstruction of the anatomical landmarks (Aubin et al., 1997), but using only one radiograph for achieving 3D reconstructions together with prior data of the subject (Boisvert et al., 2009) is an interesting and open issue of research that would allow to reduce patients' exposure to radiation.

The method proposed by Kadoury et al. (2009a) is probably the spline-based method with higher accuracy of 3D reconstruction for patients with moderate scoliosis. Still, it requires considerable computation time (2.4 min) in addition to the interaction time needed for identifying the splines. Furthermore, it was only validated for mild scoliotic patients and neither clinical indices nor vertebrae pose were assessed. This method finds an initial reconstruction using a statistical approach that is refined using image processing subject to several restrictions. Comparison with the proposed method (table 5.1) shows comparable mean reconstruction errors of the endplates for a similar sample (moderate scoliosis), while requiring much less computation time ($\sim$3 s). Additionally, the proposed method achieves lower standard deviation (0.3 vs 0.9 mm) and lower maximum error (2.3 vs 4.7 mm) for the endplates, which demonstrates more robustness locating these landmarks. Moreover, when testing the method on patients with severe scoliosis, the endplates' results remain more robust than the results presented in Kadoury et al. (2009a) for moderate scoliosis. Concerning pedicles reconstruction, the method presented here achieves higher mean reconstruction errors (3.5 vs 2.0 mm), since they are completely inferred by the statistical model with no direct clues from the operator nor from the content of the radiographic images. Nevertheless, the proposed method presents much lower standard deviation (0.4 vs 1.5 mm) and lower maximum error (4.3 vs 5.5mm), which again shows more stable results despite the average error being higher.

It is also important to mention that the method proposed by Kadoury *et. al* uses a considerably larger database for creating the statistical model (732 vs ours 295 exams), which may have direct impact on results. This was needed since the statistical approach proposed by the authors is based on Local Linear Embedding (LLE) and this technique is sensible to insufficient sampling (van der Maaten et al., 2009). In fact, despite LLE and several other dimensionality reduction techniques showing good results on artificial datasets, experiments with real-world data show that PCA often outperforms them (van der Maaten et al., 2009). Therefore, the method proposed here should be able to better modelling the population when fewer cases are available for building the statistical model, which may happen on other kinds of deformities, or in institutions without access to such amount of data.

Comparison with the results presented by Dumas et al. (2008) shows that vertebrae location accuracy was considerably improved by the method proposed here (table 5.2) while requiring less interaction (1.5 min vs 5 min). Moreover, results with severe scoliosis were comparable with moderate scoliosis, which shows the method ability for locating vertebrae even on more severe cases. In terms of vertebrae orientation, results were comparable with Dumas et al. (2008) with the exception of vertebrae rotation on Y axis that was more accurately estimated by our method. Axial rotation is considerably less accurately estimated than rotations about X and Y axes, and is sensible to an increase on the severity of scoliosis. This was expected since, unlike the rotation about X and Y axes that are calculated with the reconstructed endplates only, axial rotation is calculated using the reconstructed pedicles, which have inferior accuracy. Nevertheless, the estimation of axial rotation with the reference method has higher inter-user variability than the remaining rotations even on *in vitro* experiments (Dumas et al., 2004). These errors tend to increase on *in vivo* scenarios due to difficulties identifying pedicles, making the inter/intra-user variablity on the apical vertebrae raising up to 8° (Delorme et al., 2003). Consequently, it is difficult to properly quantify the accuracy of pedicles' reconstruction as well as axial rotation.

A proper comparison with Humbert et al. (2009) was not possible since the authors did not perform an accuracy study on neither vertebrae location, orientation, nor clinical indices. Nevertheless, the method presented here has considerably less user interaction (∼1 minute less). While both methods require users to identify the control points of one spline per view, the method proposed by Humbert et al. (2009) requires 13 additional actions from the user for each view, including manual adjustment of the shape of predefined vertebrae. Moreover, the

times presented in Humbert et al. (2009) benefit from radiographs captured with superior image quality and with no image distortion along the Z axis due the use of an EOS system (Dubousset et al., 2005) instead of standard radiographic equipment with cone-beam x-rays. These features of the EOS facilitate the identification of the splines as well as of anatomical features, which may contribute for faster times and higher accuracy.

Results of the accuracy study of clinical indices (table 5.3) show that no significant differences were found for patients with moderate scoliosis ($p \leq 0.05$). Despite there is a large variation on the orientation of the plane of maximum deformation, inter-observer precision of the reference method reaches a variability of 20.4° (Delorme et al., 2003). On the test set of severe scoliotic patients, two indices presented significant differences: the Cobb angle at the maximum plane of deformity and the angle of kyphosis. However, both these indices present acceptable mean differences for clinical practice. Additionally, the inter-observer variability of kyphosis for the reference method is of 2.8°. Therefore, we can conclude that the method presented here is suitable for clinical evaluation of both moderate and severe scoliosis (Cobb angle up to 70°).

5.8 Summary

This chapter proposed and assessed a novel 3D reconstruction method of the scoliotic spine. The method is based on two fundamental concepts: (i) the ability of articulated models for inferring missing information and (ii) the exploration of the position of the splines' control points for improving reconstruction accuracy without considerably increasing interaction time. To the best of our knowledge, these two features enabled to achieve the fastest reconstructions as well as the highest accuracy locating the vertebral bodies and the centres of their endplates. Experiments were done with both expert and non-expert users that shown that the proposed method is able to estimate clinical indices for both classes of users with no significant differences to fully manual methods. This combination of low user-interaction with reliable output makes this method suitable for use in clinical routine examinations.

CHAPTER 6

Conclusion

This thesis addressed the problem of 3D reconstruction of the spine in standard radiological environments, which involves two subproblems: geometrical calibration and recovery of the global shape of the spine. In order for this technology to become widely available to standard radiological setups, both these subproblems have to be solved with limited user input, low computation time and minimal changes to protocols, while delivering reconstructions sufficiently accurate for clinical evaluations. Such requirements were completely achieved in terms of global shape recovery of the spine where the method proposed in chapter 5 has proven to be both fast and accurate. Regarding the subproblem of geometrical calibration, it was shown that the proposed method (chapters 3 and 4) enables to recover the structure and orientation of the spine even without user-input (required for optimising parameters) and, therefore, it also fulfils the above requirements.

Other studies have shown that the calculation of angular indices of the spine was possible without requiring calibration objects (Kadoury et al., 2007a). However, to the best of our knowledge, recovering scale without calibration objects was never achieved in biplanar radiography of the spine (Cheriet et al., 1999a; Kadoury et al., 2007a). This thesis has shown that including a distance in the geometrical model that can be easily and accurately measured on site with inexpensive equipment enables to recover scale. Additionally, we have also shown that

the new model improves calibration results when calibration objects are available. For validating these findings we used validation methods where a ground truth is available, namely by generating planar data from *in vivo* CT, and by performing experiments in real environment with dried specimen. Unlike *in vivo* experiments, these methods offer a reliable ground truth that allows to accurately measure the real performance of the method. Nevertheless, future work may include *in vivo* validation for measuring the precision of the method or for comparing results with previous methods, in particular, calculating clinical indices.

Concerning the 3D reconstruction of the spine, the method proposed in this thesis allowed to decrease user-interaction, while increasing accuracy of vertebrae location. Moreover, it was shown that reliable clinical indices may be obtained rapidly by both experts and non-expert users. This is an important step in this area since methods that aim at reducing user-interaction have not yet proven their accuracy calculating these indices, which, at the end, are one of the most useful outputs for physicians. Validation was performed *in vivo* by comparing results with the gold-standard method, i.e. manual biplanar reconstructions by an expert, which allowed us to compare inter-method differences with inter/intra-observer variability and, thus, conclude that the proposed approach presents levels of accuracy within the precision of the manual method.

The next sections summarise the contributions of this thesis and point out several research directions that this thesis opens.

6.1 Main contributions

Concerning geometrical calibration of biplanar radiography, our contributions for making calibration methods easier to use on standard radiological setups were directed to minimise requirements in terms of calibration objects while providing 3D reconstructions with the correct scale. In summary, our contributions to this area were:

- A method for geometrical calibration of general biplanar radiography in standard setups. A well known calibration method that is generally used on spine radiography was extended by including a distance in the model that can be easily measured on site. We have shown that the new model coupled with an appropriate optimiser enables to improve the accuracy of 3D reconstructions of the spine. Additionally, the method remains general to be used with other anatomic structures.

- Scale recovery and accurate 3D reconstructions on biplanar radiography of the spine without calibration objects. This thesis has shown that the proposed method coupled with an appropriate measuring device enables to reconstruct the spine with correct scale without requiring calibration objects, something that was not previously achieved. This allows radiographs to be free of artifacts, less changes in the radiological protocols, and subjects do not need to wear a calibration jacket/belt, making the method easier to implement on standard radiological setups. It also makes possible to perform retrospective studies on radiographs acquired with fixed setups.

- Study of the effect of calibration object's position and size on the accuracy of 3D reconstructions of the spine. This thesis has shown that when using a small calibration object to correct scale, the accuracy of the proposed method can be improved and achieve comparable results to methods that utilise more complex objects. Only two radiopaque pellets are required, providing minimal impact to the content of radiographs. Additionally, a study was made that showed that the orientation, placement and size of calibration objects affect 3D reconstructions. This study was particular important since it has shown that while point-to-point reconstruction errors do not seem to be affected by the position of the object, the fact is that the spinal length is highly dependent of this parameter. This hidden link may explain why previous methods that used a small calibration object for scaling reconstructions show significative differences on this index when compared to large calibration apparatus (e.g. Kadoury et al. (2007b)).

Regarding the 3D reconstruction of calibrated biplanar radiographs of the spine, our efforts were directed towards delivering accuracy with minimal user interaction and, thus, making possible to use 3D reconstruction systems on routine clinical environment. In summary, our main contribution was a computational method that provides:

- The fastest reconstruction times ever. Reconstruction time is dominated by user interaction, which was decreased by one minute (from 2.5 to 1.5 minutes for mild scoliotic patients) when compared to the fastest reported times of user interaction for reconstructing the spine. In fact, the method proposed here requires less 26 actions from the user than the previously fastest method. This achievement is especially important for routine clinic where technicians have very limited time per patient and are not experts in computer tools.

- The higher accuracy ever locating vertebral bodies and endplates by a semi-

supervised method. The efficient use of the user input enabled to improve the location of vertebral bodies and their endplates, even when compared with image-based methods. Additionally, results remained equally accurate on patients with severe scoliosis showing that the method is able to handle a wide range of patients, which increases its usefulness and applicability.

- Reliable calculation of clinical indices, even by non-experts. The estimation of indices that physicians use to assess scoliosis has shown to be comparable to highly-supervised reconstruction methods. Additionally, reliable indices were also achieved by non-expert users with very limited training, showing that, for the first time, fast 3D reconstructions may be performed with useful output even when trained and experimented users are not available. This illustrates the simplicity of the input and facilitates the introduction of 3D reconstructions of the spine from biplanar radiographs on standard radiological services.

Finally, this thesis reviewed the work that was carried out on calibration and 3D reconstruction of multi-planar radiography of the spine, offering a comprehensive analysis and identifying future directions of research in these areas.

6.2 Future work

Ongoing work is related to the *in vivo* validation of the calibration method for determining the accuracy of clinical indices. Despite it was not yet possible to use the proposed calibration method together with a previously validated method for comparing results, recently, calibrated data were made available. The data set is composed by exams calibrated with a vest with several pellets (Cheriet et al., 2007). While the exact method proposed here will not be possible to validate, some of their principles may be put under test. In particular, focal distance determined by the reference method may be used as input to the proposed method.

Other issue for future work concerning the calibration method is automatising the input of the optimisation procedure by using statistical models. One interesting approach would be bundling the reconstruction approach proposed in this thesis with the calibration method. Such approach would target to minimise projection/triangulation errors and input fitting errors, while maximising statistical validity. The initial solution would be estimated as proposed in section 4.1.3, which showed to provide accurate reconstruction by itself. An alternative

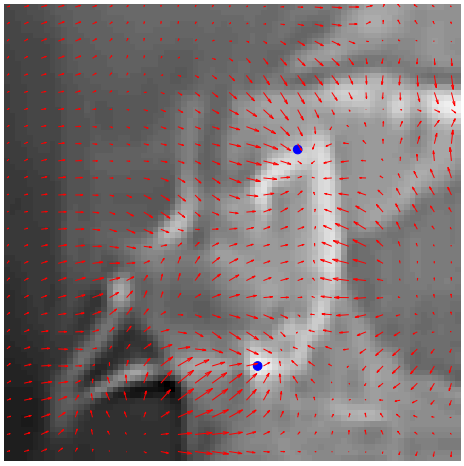


FIGURE 6.1: Posterio-anterior radiograph of a pedicle, after denoising (Dabov et al., 2007), and illustration of the Gaussian gradient of the image. (Arrows show gradient direction with size proportional to the gradient magnitude. The tips of the pedicle, manually identified by an expert, are represented by two blue dots.)

approach to provide input would be the use of a vest with calibration pellets like proposed by Cheriet et al. (2007). This approach would provide accurate and easily automatised identification of stereo-corresponding points with the disadvantage of introducing artifacts in the radiographs. This issue could be alleviated by performing a study to determine the minimum number of pellets needed for accurate calibrations using the model proposed in this thesis.

Concerning the 3D reconstruction of the spine, results from the *in vivo* validation show that endplates estimation already outperforms image-based methods, but that pedicles may be improved. Since endplates are well localised and an approximated estimation of the pedicles' position is available, we have a solid base for applying image processing techniques. A possible approach would be to explore 2D pedicles segmentation. Pedicles on the frontal view are visible for all vertebrae, except for cases of high axial rotation where one of the pedicles may be occluded, or when calibration pellets overlap pedicles. On the lateral view pedicles are usually superimposed and on the upper thoracic part of the spine they are hardly visible due to overlapping structures. Therefore, it makes sense to first segment pedicles on the frontal view and then use results to constrain their location on the lateral view. Since pedicles on the frontal view have a typical shape and there is intensity contrast with the neighbour pixels, it would be interesting to use binary classifiers, like support vector machines, to train if a given block of pixels may correspond to a pedicle tip or not. Instead of using a block of pixels directly, features descriptor should be more appropriate for this

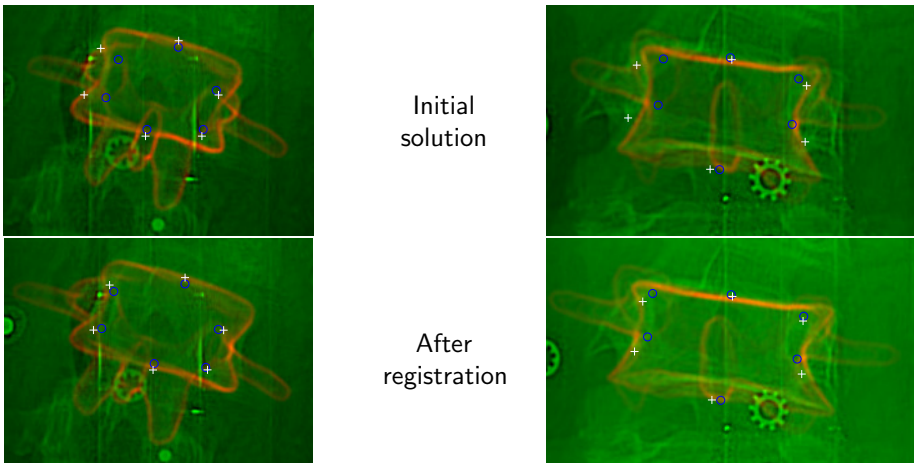


Figure 6.2: Preliminary results of 2D/3D registration: Digitally Reconstructed Radiographs (DRR) are generated for frontal and lateral views simultaneously using the GPU. Similarity metrics between DRR and radiographs are optimised for both views simultaneously. (This figure shows the projection to the frontal radiograph of the gradient of the DRR (at red) before and after registration for two vertebrae. The initial solution was given by the method proposed in chapter 5. White crosses represent 6 anatomical points identified by an expert (ground-truth) and blue circles represent the same points on the reconstruction.)

task. In particular, gradient-based descriptors such as histograms of oriented gradients (HOG) (Lowe, 2004; Dalal et al., 2005) or accumulators of gradients (Bay et al., 2006) should provide better results due to their capacity to detect singular features, and to the patterns that pedicles show in the gradient of the image (figure 6.1). Another interesting experiment would be to use probability scores of a given pixel belonging to a pedicle (Doré et al., 2007) instead of using pixels intensities directly. Since the reconstruction method proposed on chapter 5 accurately determines orientation on the frontal plane and the coordinates of the centres of the endplates, this data may be used for making the block of histograms invariant to scale and rotation. In addition, the initial reconstruction may be used for defining small regions of interest and, thus, reducing the chances of finding false positives. Two-dimensional statistical models may also be used that relate endplates positions with pedicles positions for helping to discard eventual false positives. Having the pedicles detected on the frontal radiograph, finding them on the lateral radiograph may be constrained to a search on the epipolar lines. This could be coupled with a variant of the spline-based method proposed here that would find the probable 3D location of the pedicles based on the results of the segmentation on the frontal view.

Another image-based approach that may be used for improving reconstruc-

tions and, additionally, for recovering vertebrae's shape, is 2D/3D registration of deformable 3D models of vertebrae. Other studies have tackled this problem using the projected silhouette of a deformed 3D model to fit on a binary edge map of the radiographs (Benameur et al., 2003, 2005; Kadoury et al., 2009a). We believe that better results can be achieved using as initial solution the result of the proposed 3D reconstruction method, and by performing a texture-based registration instead of only using the edges of the image. We are currently working on this approach where Digitally Reconstructed Radiographs (DRR) are generated from 3D mesh models of vertebrae (appendix B) and then registered on the radiographs by optimising a similarity metric (appendix C). For fast registrations at low budget, DRR are generated using a common Graphical Unit Processor (GPU). Currently, vertebrae deformation is controlled by the location of six 3D points (appendix A), which have shown to be insufficient to improve reconstruction results. Even though, preliminary results show that vertebrae's shape may be improved using this approach (figure 6.2). Future work will be directed to use 3D Point Distribution Models (PDM) of vertebrae reconstructed from CT for properly controlling vertebrae shape. Additionally, improving the initial 3D reconstruction with 2D segmentation before performing 2D/3D registration should increase the chances of success recovering vertebrae shape.

In the last years, research in this area was mainly directed to decrease user-interaction, and from hours of interaction we are now in less than two minutes. Therefore, it is natural that new methods start to arise in the next years that try reconstructing the spine with even less user interaction. For fulling automating the proposed method, the spine midlines would have to be extracted without user interaction. The work of Duong et al. (2010) has accomplished to extract the spine midline in postero-anterior radiographs, showing that this could be done for this view. Detecting the spine midline in the lateral view it is still an open topic of research despite Kadoury et al. (2009b) showing promising results. Other difficulties of this task are detecting the end vertebrae with accuracy. Due to the bad quality of the lateral image, first detecting the midline on the PA with the method proposed by Duong et al. (2010) and then using a hybrid approach for the lateral view using both image and statistical data relating the spline on the PA with the spline on the lateral may help solving this issue.

Finally, another path of research that would provide additional benefits to scoliotic patients would be the 3D reconstruction of the spine using only the frontal radiograph. The method proposed here is able to tackle such issue if the projection matrix is known, but preliminary experiments show that further data

is required. Using data from previous reconstructions of the same patient may be worthwhile to explore since this would expose the patient to less radiation, even if reconstructions lose some accuracy.

Bibliography

Abdel-Aziz, Y. and Karara, H. (1971). Direct linear transformation from comparator coordinates into object space coordinates in close-range photogrammetry. In *Symposium on Close-Range Photogrammetry*, pages 1–18.

Aubin, C., Dansereau, J., Parent, F., Labelle, H., and de Guise, J. A. (1997). Morphometric evaluations of personalised 3D reconstructions and geometric models of the human spine. *Medical & Biological Engineering & Computing*, 35(6):611–618.

Aubin, C., Dansereau, J., Petit, Y., Parent, F., De Guise, J., and Labelle, H. (1998). Three-dimensional measurement of wedged scoliotic vertebrae and intervertebral disks. *European Spine Journal*, 7(1):59–65.

Aubin, C.-E., Descrimes, J.-L., Dansereau, J., Skalli, W., Lavaste, F., and Labelle, H. (1995). Geometrical modelling of the spine and thorax for biomechanical analysis of scoliotic deformities using finite element method. *Annales de chirurgie*, 49(8):749–761.

Bavoil, L. and Myers, K. (2008). Order independent transparency with dual depth peeling. Technical report, NVIDIA.

Bay, H., Tuytelaars, T., and Van Gool, L. (2006). SURF: Speeded up robust features. In *European Conference on Computer Vision (ECCV)*, pages 404–417.

Bellefleur, C. (2001). Revue des logiciels pour les requis des modules d'analyse. Technical report, LIS 3D – Laboratoire d'informatique de la scoliose en trois dimensions.

Benameur, S., Mignotte, M., Labelle, H., and Guise, J. A. D. (2005). A hierarchical statistical modeling approach for the unsupervised 3-D biplanar recon-

struction of the scoliotic spine. *IEEE Transactions on Biomedical Engineering*, 52(12):2041–2057.

Benameur, S., Mignotte, M., Parent, S., Labelle, H., Skalli, W., and De Guise, J. (2001). 3D biplanar reconstruction of scoliotic vertebrae using statistical models. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 577–582.

Benameur, S., Mignotte, M., Parent, S., Labelle, H., Skalli, W., and de Guise, J. (2003). 3D/2D registration and segmentation of scoliotic vertebrae using statistical models. *Computerized Medical Imaging and Graphics*, 27(5):321–337.

Bertrand, S., Skalli, W., Delacherie, L., Bonneau, D., Kalifa, G., and Mitton, D. (2006). External and internal geometry of european adults. *Ergonomics*, 49(15):1547–1564.

Blanchard, B. and Elbaroudi, F. (2008). Computerized imaging method for a three-dimensional reconstruction from two-dimensional radiological images; implementation device. WO Patent WO/2008/012479 A1.

Boisvert, J., Cheriet, F., Pennec, X., and Ayache, N. (2009). 3D reconstruction of the human spine from radiograph(s) using a multi-body statistical model. *Progress in biomedical optics and imaging*, 10(37).

Boisvert, J., Cheriet, F., Pennec, X., Labelle, H., and Ayache, N. (2008a). Articulated spine models for 3-D reconstruction from partial radiographic data. *IEEE Transactions on Biomedical Engineering*, 55(11):2565–2574.

Boisvert, J., Cheriet, F., Pennec, X., Labelle, H., and Ayache, N. (2008b). Geometric variability of the scoliotic spine using statistics on articulated shape models. *IEEE Transactions on Medical Imaging*, 27(4):557–568.

Boisvert, J., Cheriet, F., Pennec, X., Labelle, H., and Ayache, N. (2008c). Principal deformations modes of articulated models for the analysis of 3D spine deformities. *Electronic Letters on Computer Vision and Image Analysis*, 7(4):13–31.

Bookstein, F. (1989). Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on pattern analysis and machine intelligence*, 11(6):567–585.

Bougnoux, S. (1998). From projective to euclidean space under any practical situation, acriticism of self-calibration. In *International Conference on Computer Vision*, pages 790–796.

Bouguet, J. Y. (2010). Camera calibration toolbox for Matlab. http://www.vision.caltech.edu/bouguetj/calib_doc/ (accessed on July, 2010).

Bradski, G. (2000). The OpenCV library. *Dr. Dobb's Journal of Software Tools*.

Bras, A. L., Laporte, S., Mitton, D., de Guise, J. A., and Skalli, W. (2003). Three-dimensional (3D) detailed reconstruction of human vertebrae from low-dose digital stereoradiography. *European Journal of Orthopaedic Surgery & Traumatology*, 13(2):57–62.

Canny, J. (1986). A computational approach to edge detection. *IEEE transactions on pattern analysis and machine intelligence*, 8(6):679–697.

Cassar-Pullicino, V. and Eisenstein, S. (2002). Imaging in scoliosis: What, why and how? *Clinical Radiology*, 57(7):543–562.

Cheriet, F., Dansereau, J., Petit, Y., Aubin, C., Labelle, H., and De Guise, J. A. (1999a). Towards the self-calibration of a multiview radiographic imaging system for the 3D reconstruction of the human spine and rib cage. *International Journal of Pattern Recognition and Artificial Intelligence*, 13(5):761–779.

Cheriet, F., Delorme, S., Dansereau, J., Aubin, C., De Guise, J., and Labelle, H. (1999b). Intraoperative 3D reconstruction of the scoliotic spine from radiographs. *Annales de Chirurgie*, 53(8):808–815.

Cheriet, F., Laporte, C., Kadoury, S., Labelle, H., and Dansereau, J. (2007). A novel system for the 3-D reconstruction of the human spine and rib cage from biplanar x-ray images. *IEEE Transactions on Biomedical Engineering*, 54(7):1356–1358.

Cheriet, F. and Meunier, J. (1999). Self-calibration of a biplane x-ray imaging system for an optimal three dimensional reconstruction. *Computerized Medical Imaging and Graphics*, 23(3):133–141.

Coleman, T. F. and Li, Y. (1996). An interior trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization*, 6(2):418–445.

Cootes, T., Taylor, C., Cooper, D., and Graham, J. (1992). Training models of shape from sets of examples. In *Proceedings of the British Machine Vision Conference*, pages 9–18.

Cootes, T., Taylor, C., Cooper, D., Graham, J., et al. (1995). Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59.

Corouge, I., Gouttard, S., and Gerig, G. (2004). A statistical shape model of individual fiber tracts extracted from diffusion tensor mri. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 671–679.

Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K. (2007). Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095.

Dalal, N., Triggs, B., Rhone-Alps, I., and Montbonnot, F. (2005). Histograms of oriented gradients for human detection. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 886–893.

Dansereau, J. and Stokes, I. (1988). Measurements of the three-dimensional shape of the rib cage. *Journal of Biomechanics*, 21(11):893–901.

Delorme, S., Petit, Y., de Guise, J., Labelle, H., Aubin, C.-E., and Dansereau, J. (2003). Assessment of the 3-D reconstruction and high-resolution geometrical modeling of the human skeletal trunk from 2-D radiographic images. *IEEE Transactions on Biomedical Engineering*, 50(8):989–998.

Deschenes, S. and de Guise, J. (2002). Wavelet-based automatic segmentation of the vertebral bodies in digital radiographs. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 4, pages 3868–3871.

Deschenes, S., Godbout, B., Branchaud, D., Mitton, D., Pomero, V., Bleau, A., Skalli, W., and de Guise, J. A. (2003). Three-dimensional reconstruction of the human spine from bi-planar radiographs: using multiscale wavelet analysis and spline interpolators for semi-automation. In *SPIE Medical Imaging*, pages 754–761.

Doré, V., Duong, L., Cheriet, F., and Cheriet, M. (2007). Towards segmentation of pedicles on posteroanterior x-ray views of scoliotic patients. In *Image Analysis and Recognition*, volume 4623 of *LNCS*, pages 1028–1039.

Dubousset, J., Charpak, G., Dorion, I., Skalli, W., Lavaste, F., Deguise, J., Kalifa, G., and Ferey, S. (2005). A new 2D and 3D imaging approach to musculoskeletal physiology and pathology with low-dose radiation and the standing position:

the EOS system. *Bulletin de l'Académie nationale de médecine*, 189(2):287–297.

Dumas, R., Blanchard, B., Carlier, R., de Loubresse, C., Le Huec, J., Marty, C., Moinard, M., and Vital, J. (2008). A semi-automated method using interpolation and optimisation for the 3D reconstruction of the spine from bi-planar radiography: a precision and accuracy study. *Medical and Biological Engineering and Computing*, 46(1):85–92.

Dumas, R., Le Bras, A., Champain, N., Savidan, M., Mitton, D., Kalifa, G., Steib, J. P., de Guise, J. A., and Skalli, W. (2004). Validation of the relative 3D orientation of vertebrae reconstructed by bi-planar radiography. *Medical Engineering & Physics*, 26(5):415–422.

Dumas, R., Mitton, D., Laporte, S., Dubousset, J., Steib, J. P., Lavaste, F., and Skalli, W. (2003). Explicit calibration method and specific device designed for stereoradiography. *Journal of Biomechanics*, 36(6):827–834.

Duong, L., Cheriet, F., and Labelle, H. (2010). Automatic detection of scoliotic curves in posteroanterior radiographs. *IEEE Transactions on Biomedical Engineering*, 57(5):1143–1151.

Eggert, D., Lorusso, A., and Fisher, R. (1997). Estimating 3-D rigid body transformations: a comparison of four major algorithms. *Machine Vision and Applications*, 9(5):272–290.

Everitt, C. (2001). Interactive order-independent transparency. Technical report, NVIDIA.

Fairbank, J. (2002). Idiopathic scoliosis. In Bulstrode, C., Buckwalter, J., Carr, A., Marsh, L., Fairbank, J., Wilson-MacDonald, J., and Bowden, G., editors, *Oxford Textbook of Orthopedics and Trauma*, chapter 2.1.4.2, pages 574–582. Oxford University Press, 1st edition.

Faugeras, O., Luong, Q., and Maybank, S. (1992). Camera self-calibration: Theory and experiments. In *European Conference on Computer Vision*, pages 321–334.

Forsyth, D. and Ponce, J. (2003). *Computer vision: a modern approach*. Prentice Hall, New Jersey.

Fritsch, F. and Carlson, R. (1980). Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, 17(2):238–246.

Gille, O., Champain, N., Benchikh-El-Fegoun, A., Vital, J.-M., and Skalli, W. (2007). Reliability of 3D reconstruction of the spine of mild scoliotic patients. *Spine*, 32(5):568–573.

Greenspan, A. (2004). *Orthopedic Imaging: A Practical Approach*. Lippincott Williams & Wilkins, 4th edition.

Hartley, R. (1994). Euclidean reconstruction from uncalibrated views. *Applications of invariance in computer vision*, pages 235–256.

Hartley, R. and Sturm, P. (1997). Triangulation. *Computer Vision and Image Understanding*, 68(2):146–157.

Hartley, R. and Zisserman, A. (2003). *Multiple view geometry in computer vision*. Cambridge University Press New York, NY, USA, 2nd edition.

Horn, B. et al. (1987). Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4(4):629–642.

Hornegger, J. and Tomasi, C. (1999). Representation issues in the ml estimation of camera motion. In *International Conference on Computer Vision (ICCV)*, pages 640–647.

Humbert, L., De Guise, J., Godbout, B., Parent, S., Dubousset, J., and Skalli, W. (2007). 3D reconstruction of the spine from biplanar x-rays using longitudinal and transversal inferences. *International Journal of Computer Assisted Radiology and Surgery*, 2(Sup 1):S452.

Humbert, L., Guise, J. D., Aubert, B., Godbout, B., and Skalli, W. (2009). 3D reconstruction of the spine from biplanar x-rays using parametric models based on transversal and longitudinal inferences. *Medical Engineering & Physics*, 31(6):681–687.

Ibanez, L., Schroeder, W., Ng, L., and Cates, J. (2005). *The ITK Software Guide*. Kitware, http://www.itk.org/ItkSoftwareGuide.pdf, 2nd edition.

Jolliffe, I. (2002). *Principal component analysis*. Springer-Verlag, New York, 2nd edition.

Kadoury, S. and Cheriet, F. (2006). X-ray restoration with adaptive pde filter for an accurate calibration of the acquisition system for the 3D reconstruction of the human spine. *Computer Assisted Radiology and Surgery*, 1(Sup 7):470.

Kadoury, S., Cheriet, F., Dansereau, J., and Labelle, H. (2007a). Three-dimensional reconstruction of the scoliotic spine and pelvis from uncalibrated biplanar x-ray images. *Journal of Spinal Disorders & Techniques*, 20(2):160–167.

Kadoury, S., Cheriet, F., and Labelle, H. (2009a). Personalized x-ray 3D reconstruction of the scoliotic spine from hybrid statistical and image-based models. *IEEE Transactions on Medical Imaging*, 28(9):1422–1435.

Kadoury, S., Cheriet, F., and Labelle, H. (2009b). Segmentation of scoliotic spine silhouettes from enhanced biplanar x-rays using a prior knowledge bayesian framework. In *IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI)*, pages 478–481.

Kadoury, S., Cheriet, F., and Labelle, H. (2010). Self-calibration of biplanar radiographic images through geometric spine shape descriptors. *IEEE Transactions on Biomedical Engineering*, 57(7):1663–1675.

Kadoury, S., Cheriet, F., Laporte, C., and Labelle, H. (2007b). A versatile 3D reconstruction system of the spine and pelvis for clinical assessment of spinal deformities. *Medical & Biological Engineering & Computing*, 45(6):591–602.

Karachalios, T., Sofianos, J., Roidis, N., Sapkas, G., Korres, D., and Nikolopoulos, K. (1999). Ten-year follow-up evaluation of a school screening program for scoliosis: Is the forward-bending test an accurate diagnostic criterion for the screening of scoliosis? *Spine*, 24(22):2318–2324.

Labelle, H., Bellefleur, C., Joncas, J., Aubin, C., and Cheriet, F. (2007). Preliminary evaluation of a computer-assisted tool for the design and adjustment of braces in idiopathic scoliosis: a prospective and randomized study. *Spine*, 32(8):835–843.

Labelle, H., Dansereau, J., Bellefleur, C., and Jéquier, J. (1995a). Variability of geometric measurements from three-dimensional reconstructions of scoliotic spines and rib cages. *European Spine Journal*, 4(2):88–94.

Labelle, H., Dansereau, J., Bellefleur, C., and Poitras, B. (1996). Three-dimensional effect of the boston brace on the thoracic spine and rib cage. *Spine*, 21(1):59–64.

Labelle, H. et al. (1995b). Comparison between preoperative and postoperative three-dimensional reconstructions of idiopathic scoliosis with the cotrel-dubousset procedure. *Spine*, 20(23):2487–2492.

Lemieux, L., Jagoe, R., Fish, D., Kitchen, N., and Thomas, D. (1994). A patient-to-computed-tomography image registration method based on digitally reconstructed radiographs. *Medical physics*, 21:1749–1760.

Lepetit, V. and Fua, P. (2005). Monocular model-based 3D tracking of rigid objects: A survey. *Foundations and Trends in Computer Graphics and Vision*, 1(1):1–89.

Levy, A., Goldberg, M., Mayo, N., Hanley, J., and Poitras, B. (1996). Reducing the lifetime risk of cancer from spinal radiographs among people with adolescent idiopathic scoliosis. *Spine*, 21(13):1540–1547.

Lewis, R. and Torczon, V. (2002). A globally convergent augmented lagrangian pattern search algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Optimization*, 12(4):1075–1089.

Lourakis, M. and Argyros, A. (2005). Is Levenberg-Marquardt the most efficient optimization algorithm for implementing bundle adjustment? In *IEEE International Conference on Computer Vision*, volume 2, pages 1526–1531.

Lowe, D. (2004). Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110.

Maesschalck, R. D., Jouan-Rimbaud, D., and Massart, D. L. (2000). The mahalanobis distance. *Chemometrics and Intelligent Laboratory Systems*, 50(1):1–18.

Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11(2):431–441.

Mattes, D., Haynor, D., Vesselle, H., Lewellyn, T., and Eubank, W. (2001). Nonrigid multimodality image registration. In *SPIE Medical Imaging: Image Processing*, pages 1609–1620.

Mitton, D., Landry, C., Veron, S., Skalli, W., Lavaste, F., and De Guise, J. (2000). 3D reconstruction method from biplanar radiography using non-stereocorresponding points and elastic deformable meshes. *Medical & Biological Engineering & Computing*, 38(2):133–139.

Mitulescu, A., Semaan, I., De Guise, J., Leborgne, P., Adamsbaum, C., and Skalli, W. (2001). Validation of the non-stereo corresponding points stereoradiographic 3D reconstruction technique. *Medical & Biological Engineering & Computing*, 39(2):152–158.

Mitulescu, A., Skalli, W., Mitton, D., and De Guise, J. (2002). Three-dimensional surface rendering reconstruction of scoliotic vertebrae using a non stereo-corresponding points technique. *European Spine Journal*, 11(4):344–352.

Moura, D. C., Barbosa, J. G., Reis, A. M., and Tavares, J. M. R. S. (2010a). A flexible approach for the calibration of biplanar radiography of the spine on conventional radiological systems. *Computer Modelling in Engineering and Sciences*, 60(2):115–138.

Moura, D. C., Barbosa, J. G., Tavares, J. M. R. S., and Reis, A. M. (2008a). Calibration of bi-planar radiography with a minimal-size calibration object. *International Journal of Computer Assisted Radiology and Surgery*, 3(Sup 1):S352–S353.

Moura, D. C., Barbosa, J. G., Tavares, J. M. R. S., and Reis, A. M. (2008b). Calibration of bi-planar radiography with a rangefinder and a small calibration object. In *Advances in Visual Computing*, volume 5358 of *LNCS*, pages 572–581.

Moura, D. C., Boisvert, J., Barbosa, J. G., and Tavares, J. M. R. S. (2009). Fast 3D reconstruction of the spine using user-defined splines and a statistical articulated model. In *Advances in Visual Computing*, volume 5875 of *LNCS*, pages 586–595.

Moura, D. C., Boisvert, J., Barbosa, J. G., Tavares, J. M. R. S., and Labelle, H. (2010b). Fast 3D reconstruction of the spine by non-expert users using a statistical articulated model. In *Research into Spinal Deformities 7*, volume 158 of *Studies in Health Technology and Informatics*, pages 268–269.

Munbodh, R., Jaffray, D., Moseley, D., Chen, Z., Knisely, J., Cathier, P., and Duncan, J. (2006). Automated 2D-3D registration of a radiograph and a cone beam ct using line-segment enhancement. *Medical physics*, 33:1398–1411.

Novosad, J., Cheriet, F., Petit, Y., and Labelle, H. (2004). Three-dimensional (3-D) reconstruction of the spine from a single x-ray image and prior vertebra models. *IEEE Transactions on Biomedical Engineering*, 51(9):1268–1639.

Papin, P., Labelle, H., Delorme, S., Aubin, C., De Guise, J., and Dansereau, J. (1999). Long-term three-dimensional changes of the spine after posterior spinal instrumentation and fusion in adolescent idiopathic scoliosis. *European Spine Journal*, 8(1):16–21.

Pearcy, M. J. (1985). Stereo radiography of lumbar spine motion. *Acta Orthopaedica*, 56(S212):1–45.

Pennec, X. (2006). Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154.

Penney, G., Batchelor, P., Hill, D., Hawkes, D., and Weese, J. (2001). Validation of a two- to three-dimensional registration algorithm for aligning preoperative ct images and intraoperative fluoroscopy images. *Medical Physics*, 28:1024–1032.

Penney, G., Weese, J., Little, J., Desmedt, P., Hill, D., and Hawkes, D. (1998). A comparison of similarity measures for use in 2D-3D medical image registration. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 1153–1161.

Pluim, J., Maintz, J., and Viergever, M. (2000). Image registration by maximization of combined mutual information and gradient information. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 103–129.

Pollefeys, M., Koch, R., and Gool, L. (1999). Self-calibration and metric reconstruction inspite of varying and unknown intrinsic camera parameters. *International Journal of Computer Vision*, 32(1):7–25.

Pomero, V., Mitton, D., Laporte, S., de Guise, J. A., and Skalli, W. (2004). Fast accurate stereoradiographic 3D-reconstruction of the spine using a combined geometric and statistic model. *Clinical Biomechanics*, 19(3):240–247.

Rajamani, K., Joshi, S., and Styner, M. (2004). Bone model morphing for enhanced surgical visualization. In *IEEE International Symposium on Biomedical Imaging: Nano to Macro (ISBI)*, pages 1255–1258.

Rogers, L. F. (1998). Introduction to skeletal radiology and bone growth. In Juhl, J. H., Crummy, A. B., and Kuhlman, J. E., editors, *Paul & Juhl's Essentials of Radiologic Imaging*, chapter 1, pages 20–28. Lippincott Williams & Wilkins, 7th edition.

Rowe, D., Bernstein, S., Riddick, M., Adler, F., Emans, J., and Gardner-Bonneau, D. (1997). A meta-analysis of the efficacy of non-operative treatments for idiopathic scoliosis. *The Journal of Bone and Joint Surgery*, 79(5):664–674.

Roweis, S. and Saul, L. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326.

Schmidt, J. and Niemann, H. (2001). Using quaternions for parametrizing 3-D rotations in unconstrained nonlinear optimization. In *Vision, Modeling, and Visualization*, pages 399–406.

Soh, L., Matas, J., and Kittler, J. (1997). Robust recognition of calibration charts. In *IEE Conference Publications*, pages 487–491.

Stokes, I. (1989). Axial rotation component of thoracic scoliosis. *Journal of Orthopaedic Research*, 7(5):702–708.

Stokes, I. A. (1994). Three-dimensional terminology of spinal deformity. a report presented to the scoliosis research society by the scoliosis research society working group on 3-D terminology of spinal deformity. *Spine*, 19(2):236–248.

Styner, M., Lieberman, J., and Gerig, G. (2003). Boundary and medial shape analysis of the hippocampus in schizophrenia. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 464–471.

The, B. (2006). *Digital radiographic preoperative planning and postoperative monitoring of total hip replacements: techniques, validation and implementation.* PhD thesis, University Medical Centre Groningen, the Netherlands.

Triggs, B. (1997). Autocalibration and the absolute quadric. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 609–614.

Trochu, F. (1993). A contouring program based on dual kriging interpolation. *Engineering with Computers*, 9(3):160–177.

Vaiton, M., Dansereau, J., Grimard, G., Beauséjour, M., Labelle, H., and de Guise, J. (2004). Évaluation d'une méthode clinique d'acquisition rapide de la géométrie 3D de colonnes vertébrales scoliotiques. *ITBM-RBM*, 25(3):150–162.

Valenza, T. (2009). Rsna 2009 blog – imaging economics. Internet. http://www.imagingeconomics.com/rsna09blog/ (accessed on July, 2010).

van der Maaten, L., Postma, E., and van den Herik, J. (2009). Dimensionality reduction: A comparative review. *Journal of Machine Learning Research*, 10:1–41.

Vasconcelos, M. J. M. and Tavares, J. M. R. S. (2008). Methods to automatically build point distribution models for objects like hand palms and faces represented in images. *Computer Modeling in Engineering & Sciences*, 36(3):213–241.

Velezis, M., Sturm, P., and Cobey, J. (2002). Scoliosis screening revisited: findings from the district of columbia. *Journal of Pediatric Orthopaedics*, 22(6):788–791.

Vidal, F. P., Garnier, M., Freud, N., Létang, J. M., and John, N. W. (2009). Simulation of x-ray attenuation on the gpu. In *Proceeding of Theory and Practice of Computer Graphics (TCPG)*, pages 25–32.

Villemure, I., Aubin, C., Grimard, G., Dansereau, J., and Labelle, H. (2001). Progression of vertebral and spinal three-dimensional deformities in adolescent idiopathic scoliosis: a longitudinal study. *Spine*, 26(20):2244–2250.

Vince, J. (2006). *Mathematics for computer graphics*. Springer Verlag, New York.

Viola, P. and Wells III, W. (1997). Alignment by maximization of mutual information. *International journal of computer vision*, 24(2):137–154.

Wang, Z., Wu, W., Xu, X., and Xue, D. (2007). Recognition and location of the internal corners of planar checkerboard calibration pattern image. *Applied Mathematics and Computation*, 185(2):894–906.

Weinstein, S., Dolan, L., Cheng, J., Danielsson, A., and Morcuende, J. (2008). Adolescent idiopathic scoliosis. *Lancet*, 371(9623):1527–1537.

Wohler, C. (2009). *3D Computer Vision: Efficient Methods and Applications*. Springer-Verlag.

Wood, G. and Marshall, R. (1986). The accuracy of dlt extrapolation in three-dimensional film analysis. *Journal of Biomechanics*, 19(9):781–785.

Working Group on 3-D Classification, T. C. (2000). Revised glossary of terms of the scoliosis research society. Internet. https://www.srs.org/professionals/glossary/glossary.asp (accessed on October, 2007).

Yazici, M., Acaroglu, E., Alanay, A., Deviren, V., Cila, A., and Surat, A. (2001). Measurement of vertebral rotation in standing versus supine position in adolescent idiopathic scoliosis. *Journal of Pediatric Orthopaedics*, 21(2):252–256.

Zhang, Z. (2000). A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence*, 22(11):1330–1334.

Zografos, V. (2009). Comparison of optimisation algorithms for deformable template matching. In *Advances in Visual Computing*, volume 5876 of *LNCS*, pages 1097–1108.

Appendices

APPENDIX A

Deformation of vertebrae models using anatomical landmarks

This appendix describes the deformation of 3D models of vertebrae based on the 6 anatomical points that are used on the methods proposed in this thesis. This technique is used for illustrating 3D reconstructions of the spine as well as for the 2D/3D registration method described in appendix C.

For generating a vertebra model from a set of anatomical landmarks (target landmarks), a generic model is deformed towards fitting the target landmarks. Generic models of vertebrae are represented by polygon meshes extracted from Computer Tomography (CT) images (Aubin et al., 1997). One model is extracted per vertebral level (T1–L5) and six anatomical landmarks are identified (source landmarks), namely the centres of the superior and inferior endplates, and the superior and inferior extremities of the two pedicles (figure A.1). Deformation is achieved using the Thin-Plate Spline transform (TPS) (Bookstein, 1989). TPS is a non-linear landmark transform that maps a set of landmarks (source landmarks) to another set of landmarks (target landmarks). Mapping is achieved trough interpolation and aims at minimizing a bending energy for ensuring that the deformation is smooth and without discontinuities.

TPS allows to compute the transform that maps the 6 anatomical landmarks of the generic 3D model of a given vertebral level to the 6 six landmarks reconstructed from a biplanar radiographic exam of any given patient. This transform is then applied to all points that make the generic vertebral model to deform it. Figure A.2 shows the result of applying the deformation algorithm to each vertebra of a reconstruction of a spine. Deformation is performed using the implementation of the Insight Toolkit (ITK) (Ibanez et al., 2005) of TPS, namely the class `ThinPlateR2LogRSplineKernelTransform`.

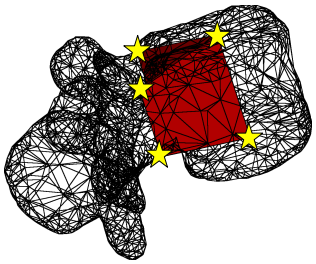


Figure A.1: Generic 3D model of a vertebra. The model is represented by a polygon mesh. Five of the six anatomical landmarks are identified in the figure with stars. For illustration purposes, the pedicles' landmarks are connected to the endplates' landmarks by two quadrilaterals.

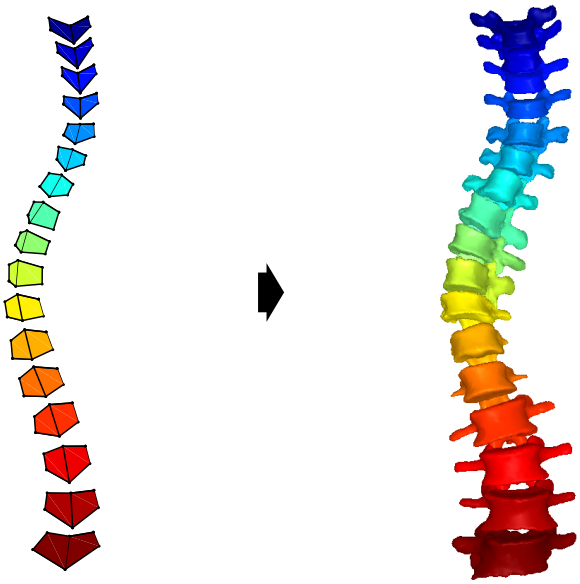


Figure A.2: Representation of the spine using 6 points per vertebra (left), and representation of the same spine using polygon meshes for describing each vertebra (right). For each vertebral level (T1–L5), a generic vertebra was deformed with the TPS transform, using the 6 anatomical points for controlling the deformation.

APPENDIX B

Generation of digitally reconstructed radiographs from polygon meshes

This appendix describes the generation of Digitally Reconstructed Radiographs (DRR) from polygon meshes of vertebrae. This technique is used during this dissertation for illustrating projections to the radiographs of the 3D reconstructions as well as for the 2D/3D registration method described in appendix C.

B.1 Brief description of the algorithm

Digitally Reconstructed Radiographs (DRR) are simulations of radiographs constructed from volumetric (3D) data. Typically, DRRs are computed from 3D exams (e.g. CT – Computer Tomography) for rigid 2D/3D registration with planar radiographs (e.g. Penney et al. (2001); Munbodh et al. (2006)). Since one of our goals for DRR generation is to enable 2D/3D registration of deformable models, both deformation and DRR synthesis have to be fast. Therefore, and as suggested in Vidal et al. (2009), vertebrae models are represented by poly-

gon meshes, and DRRs are computed from polygon meshes, instead of dense volumetric data, using the Graphic Processor Unit (GPU).

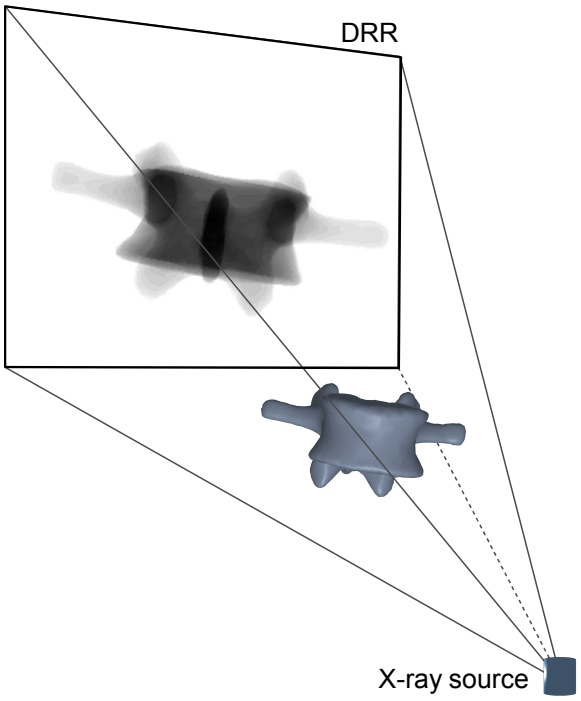


Figure B.1: Formation of a Digitally Reconstructed Radiograph (DRR) from a polygon mesh.

The algorithm that was employed here for the generation of DRRs is similar to the one proposed in Vidal et al. (2009). Basically, it simulates x-rays coming from the x-ray source to each of the pixels of the DRR for calculating their intensity (figure B.1). For a single pixel in the DRR, its intensity (I) is calculated using the attenuation law, which may be formulated as:

$$I = \exp\left(-\sum_{i=1}^{N} \mu(i) L_p(i)\right), \tag{B.1}$$

where N is the number of anatomical structures that the path of the x-ray crosses when coming from the x-ray source to the DRR's pixel, $\mu(i)$ is the linear attenuation coefficient of structure i, and $L_p(i)$ is the length of the path of the x-ray that crossed structure i. Parameter $\mu(i)$ is related to the type of material of the anatomical structure and, since only bone tissue is modelled by the mesh, a constant value was used.

One of the main differences between the algorithm implemented here and the one proposed by Vidal et al. (2009) is the calculation of $L_p(i)$. Here, a simpler

approach was used that is based on multi-pass depth-peeling implemented with the OpenGL Shading Language (GLSL) (Everitt, 2001; Bavoil and Myers, 2008). In Vidal et al. (2009) a single-pass approach was used for calculating $L_p(i)$ that creates artifacts on the DRR. These artifacts are unwanted and require the use of an extra step for detecting and correcting them by averaging the pixels in the neighbourhood of the artifacts. Even though, the method proposed by Vidal et al. (2009) should be more computationally efficient.

In addition to the DRR calculation, an optional final pass was implemented for calculating the gradient of the DRR in each of its two dimensions, as well as its magnitude. Since all these images are grayscale, the 4 color channels (RGBA) were used to store all results in a single buffer, i.e. the DRR was stored on the RED channel, the gradient magnitude was stored on the GREEN channel, the gradient in the X axis was stored on the BLUE channel, and the gradient in the Y axis was stored on the ALPHA channel. Additionally, scissoring was employed for limiting the computation area to the region of interest of the projected 3D model.

B.2 Results

For evaluating the performance of the algorithm, 50.000 DRRs were generated on a Desktop PC with an AMD Phenom II X2 550 3.10 GHz processor, 2 GBytes of RAM, equipped with a NVIDIA GeForce GTS 250 graphic card. A 3D model of a L3 vertebra composed by 1552 triangles was projected alternately to Posterio-anterior (PA) and Lateral (LAT) planes, generating DRRs with 532×276 pixels and 202×132 pixels respectively. Results with different options of the algorithm are presented on table B.1 showing that it is possible to generate more that 1000 DRRs per second with an entry-level graphic card ($\sim$110 euros, Portugal, 2010), even when copying the DRR to CPU memory for posterior calculations. Despite these numbers already allow to use DRR generation in optimisation processes, like 2D/3D registration, we are currently working on an implementation on CUDA for increasing performance even further. Figure B.2 shows the generated DRRs for the PA and LAT planes, as well as the magnitude of the gradient of the DRR. An illustration of a DRR of a spine (T1–L5) is also included for comparison with the original radiograph (figure B.3).

Table B.1: Peformance of the DRR generator implemented in OpenGL GLSL running in a NVIDIA GeForce GTS 250 graphic card. The values on the table represent the number of DRRs generated per second.

Gradient calculation	Copy to CPU memory	
	Yes	No
Yes	1125	1534
No	1311	1570

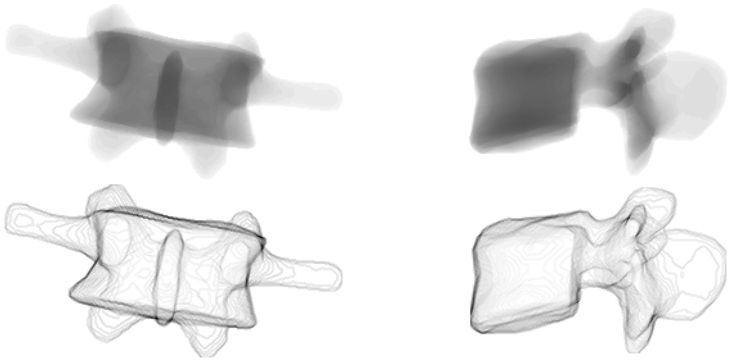
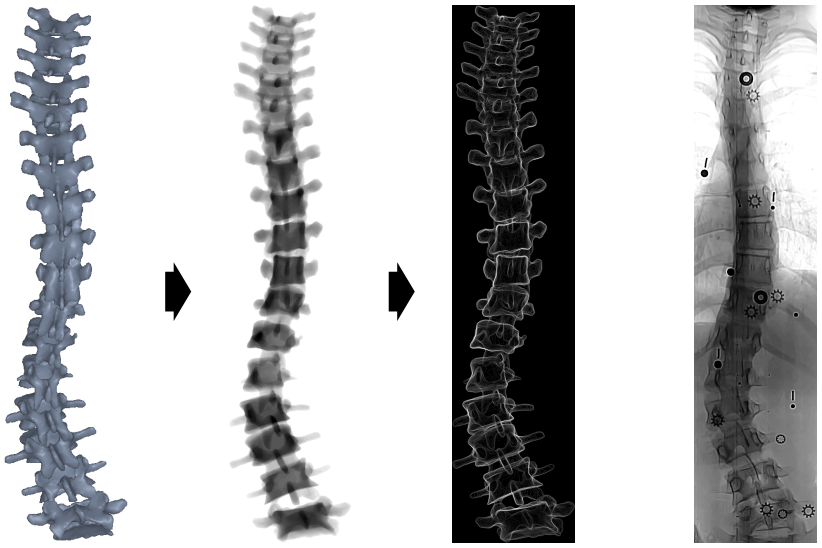


Figure B.2: Digitally reconstructed radiographs (DRR) obtained by the proposed algorithm for the Postero-anterior (at left) and Lateral (at right) planes. The top row shows the original DRR and the bottom row shows the gradient image of the DRRs (intensity inverted).



(a) (b) (c) (d)

Figure B.3: DRR (b) and respective gradient (c) of a mesh model of a spine (a), and comparison with the original posterio-anterior radiograph (d). The mesh model (a) was generated as described in appendix A from a manual reconstruction of the 6 points per vertebra by an expert. Images (c) and (d) were enhanced for visualisation purposes.

APPENDIX C

2D/3D registration of vertebrae models on radiographs

This appendix describes ongoing work for improving vertebrae pose and shape using 2D/3D registration. It starts with a brief description of the registration method and then presents preliminary results.

C.1 Brief description of the method

The method described in chapter 5 is used for estimating the position of 6 anatomical points per vertebra. For each vertebra, the 6 points are used to find an initial solution of the shape of the vertebra by deforming a 3D generic model using the Thin-Plate Spline transform (appendix A). Then, vertebrae position, orientation and scale of the model is optimised by maximising the similarity between frontal and lateral DRRs generated from the model (appendix B) and the correspondent radiographic images.

Several metrics were experimented for quantifying the similarity between DRRs and radiographs, including Normalised Cross Correlation (Lemieux et al.,

1994), Gradient Difference (Penney et al., 1998), and Mutual Information (Viola and Wells III, 1997; Mattes et al., 2001). The metrics were first assessed in 2D-2D registrations of a single DRR in a radiograph. The best results were found for a metric inspired in the gradient metric proposed in Pluim et al. (2000), which we will call Directional Gradient Difference (DGD). This metric calculates the differences between the orientations of the gradient of the DRR and radiograph, and weights these differences with the magnitude of the gradients. The gradient vectors are first determined by calculating the partial derivatives of the images with a Gaussian kernel. Let $g = \{g_x, g_y\}$ be the vector defining the gradient of a given pixel of the DRR, and $g' = \{g'_x, g'_y\}$ be the gradient vector for the corresponding pixel in the radiograph, the angle $\alpha_{g,g'}$ between the gradient vectors of the two images is defined by:

$$\alpha_{g,g'} = \arccos\left(\frac{g.g'}{|g||g'|}\right). \tag{C.1}$$

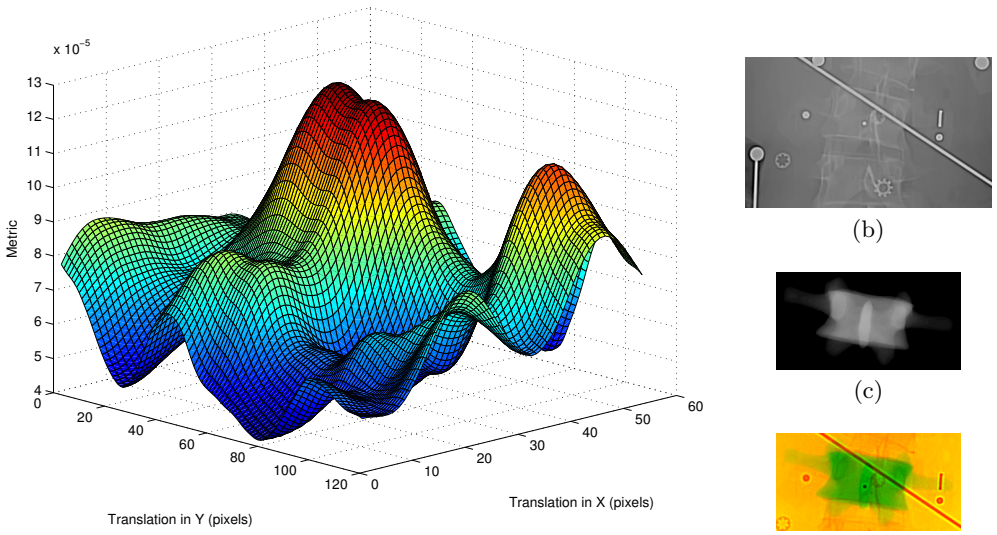
Since there is no guarantee that the direction of the gradient is the same in the two images, a weighting function w is used that enables to compare the orientation of the gradient independently of its direction:

$$w(\alpha) = \frac{\cos(2\alpha) + 1}{2}. \tag{C.2}$$

Function w has maximum value of 1 when the gradients are parallel (i.e. $\alpha = 0 \vee \alpha = \pi$), and has minimum value of 0 when the gradients are perpendicular. The original formulation proposes using as metric $w(\alpha) \times \min(g, g')$. Here, we propose defining the similarity measure m for a given pixel as:

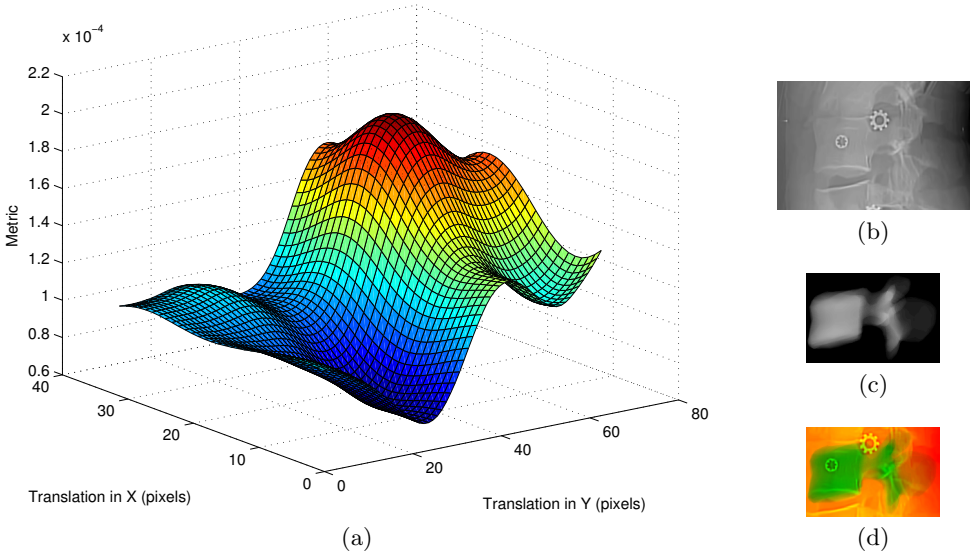
$$m = w(\alpha)^2 \times g \times g'. \tag{C.3}$$

This formulation enables to include the magnitudes of the two images, even if they have different scales, while it penalises large differences of the orientation of the gradient. If both g and g' are large, which indicates an edge or a structure, and their orientation is close then the similarity will be large, but if any of the gradients is small or the difference between orientations is large, then the similarity will be low. Finally, the metric for the intersection between the DRR and the radiograph is obtained by summing the similarities between all the pixels of the DRR and the corresponding pixels of the radiograph. For increasing computational efficiency, only pixels of the DRR with gradient magnitude above a given threshold β are considered. Figure C.1 and C.2 show the response of this metric when translating a static DRR over a posterio-anterior and lateral radiographs respectively.



(a) (b) (c) (d)

FIGURE C.1: Response of the selected metric (a) to 2D translation of the DRR (c) over the PA radiograph (b). Higher values mean higher similarity. Figure (d) shows the DRR (green) over the radiograph (orange) at the point of maximum similarity.



(a) (b) (c) (d)

FIGURE C.2: Response of the selected metric (a) to 2D translation of the DRR (c) over the LAT radiograph (b). Higher values mean higher similarity. Figure (d) shows the DRR (green) over the radiograph (orange) at the point of maximum similarity.

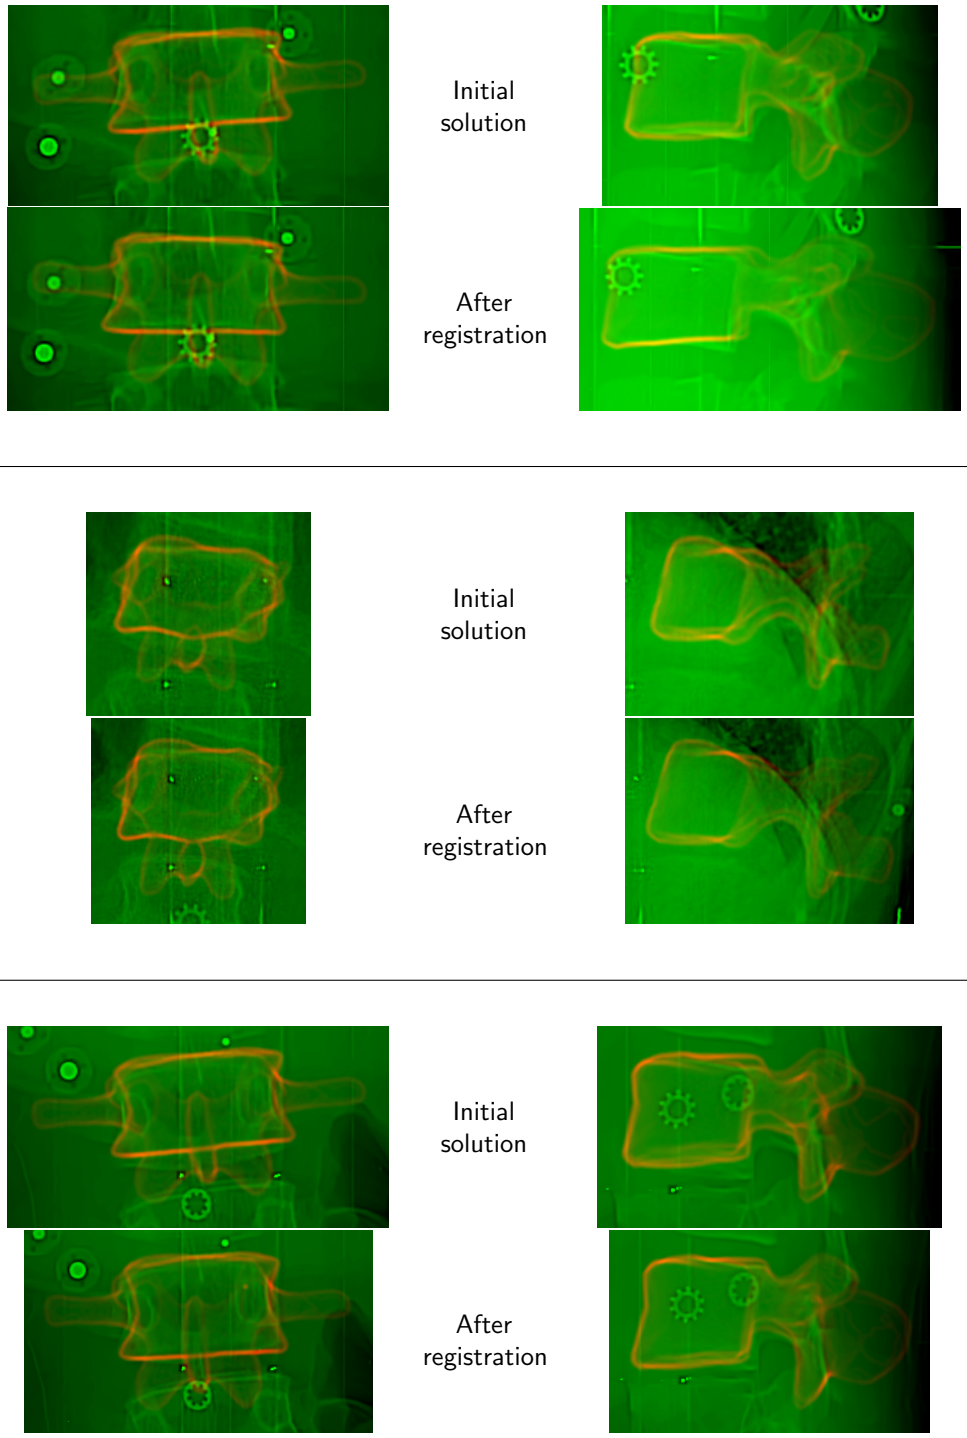


FIGURE C.3: Preliminary results of 2D/3D registration. This figure shows the DRRs' gradient (red) for the PA and LAT views overlapping the original radiographs (green), before and after registration for 3 different vertebrae.

C.2 Preliminary results

Experiments were done where the 3D position, orientation and scale of vertebrae were optimised by maximising the sum of the DGD metric of both posterio-anterior and lateral views. Optimisation was performed with a pattern search algorithm (Lewis and Torczon, 2002), which was shown to be a good optimiser for registration problems when a good initial solution is available (Zografos, 2009). Preliminary results (figure C.3) indicate that 2D/3D registration is potentially helpful for recovering the shape of vertebrae. Generally, the contours of the projected model fit better on the radiographs that in the initial solution. However, updating the 6 anatomical points based on the results of the registration process often increases errors when compared with manual identification. Future work may include (i) adding to the metric a statistical term indicating the probability of the current solution being valid, (ii) refining the 2D location of the 6 points on the radiographs with 2D image processing before running the 2D-3D registration for restricting the search space of solutions, and (iii) using 3D Point Distribution Models (PDM) (Cootes et al., 1995) for generating valid vertebrae shapes while constraining the search to the principal deformation modes.