

Registration of Non-rigidly Deforming Objects



Roberto M. Dyke
School of Computer Science & Informatics
Cardiff University

A thesis submitted in partial fulfilment of the requirement for the
degree of

Doctor of Philosophy

January 2021

Acknowledgements

Firstly, thank you to Yu-Kun Lai for giving me the opportunity to undertake my first steps into research through the *Cardiff Undergraduate Research Opportunities Programme*, and for the consistent guidance throughout my studies. Your support has allowed me to develop as a researcher.

I would also like to express my gratitude to Paul Rosin for his contributions to the project with particularly astute thought. The discussions and ideas that occurred during our meetings were consistently intriguing. It has been a great privilege to study under both your guidance.

I would like to thank my family and my fiancée for supporting me through the write-up of my thesis. Thank you to Elizabeth for making sure I had enough fun and took lots of breaks! Thank you to Kai Hormann for your gratuity, helping facilitate the final stages of producing this document. I couldn't have completed this without the support of all of you.

Further gratitude is warranted for those who helped contribute towards this work. Thank you Seana Dykes, Beryl Noë, Emma Dyke, Caleb Stride, and Zhou Feng for their assistance producing the datasets. Additionally, thank you Emma Dyke, Julian Dyke, Seana Dykes, Lauren Hudson, and Stefano Zappalà for your valuable feedback from proofreading this thesis.

I am deeply grateful to the ICT and Computing teachers from LVS Ascot (Mrs S. Featherstone-Clark, Mr J. Franklin, and Mrs K. Southard) who grew my passion for Computer Science, encouraging and inspiring me.

Finally, I would like to acknowledge my fellow postgraduate research students and office no. C/2.13 for their patience, motivation and friendship. Together, you made this one of the most wonderful times of my life, and working by you has felt more like being part of a family than a workplace.

I wish you all the very best.

This work has been supported by the Cardiff University EPSRC Doctoral Training Partnership [grant ref. EP/N509449/1].

Publications

Listed below are the publications resulting from the work undertaken for this thesis:

R. M. Dyke, Y.-K. Lai, P. L. Rosin, and G. K. L. Tam. Non-rigid registration under anisotropic deformations. *Computer Aided Geometric Design*, 71:142–156, May 2019. ISSN 0167-8396. doi:[10.1016/j.cagd.2019.04.014](https://doi.org/10.1016/j.cagd.2019.04.014)

R. M. Dyke, C. Stride, Y.-K. Lai, P. L. Rosin, M. Aubry, A. Boyarski, A. M. Bronstein, M. M. Bronstein, D. Cremers, M. Fisher, T. Groueix, D. Guo, V. G. Kim, R. Kimmel, Z. Löhner, K. Li, O. Litany, T. Remez, E. Rodolà, B. C. Russell, Y. Sahillioğlu, R. Slossberg, G. K. L. Tam, M. Vestner, Z. Wu, and J. Yang. Shape correspondence with isometric and non-isometric deformations. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '19, pages 111–119. The Eurographics Association, May 2019. ISBN 978-3-03868-077-2. doi:[10.2312/3dor.20191069](https://doi.org/10.2312/3dor.20191069)

R. M. Dyke, Z. Feng, Y.-K. Lai, P. L. Rosin, D. Guo, K. Li, R. Marin, and J. Yang. Shape correspondence of physically-based deformations. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '20. The Eurographics Association, 2020. doi:[10.2312/3dor.20201161](https://doi.org/10.2312/3dor.20201161)

R. M. Dyke, Y.-K. Lai, P. L. Rosin, S. Zappalà, S. L. Dykes, D. Guo, K. Li, R. Marin, S. Melzi, and J. Yang. Shape correspondence with non-isometric deformations. *Computers & Graphics*, 2020. ISSN 0097-8493. doi:[10.1016/j.cag.2020.08.008](https://doi.org/10.1016/j.cag.2020.08.008)

Summary

This thesis investigates the current state-of-the-art in registration of non-rigidly deforming shapes. In particular, the problem of non-isometry is considered. First, a method to address locally anisotropic deformation is proposed. The subsequent evaluation of this method highlights a lack of resources for evaluating such methods. Three novel registration/shape correspondence benchmark datasets are developed for assessing different aspects of non-rigid deformation. Deficiencies in current evaluative measures are identified, leading to the development of a new performance measure that effectively communicates the density and distribution of correspondences.

Finally, the problem of transferring skull orbit labels between scans is examined on a database of unlabelled skulls. A novel pipeline that mitigates errors caused by coarse representations is proposed.

Category

Computer Graphics · Computational Geometry · Non-Rigid Registration

Keywords

Non-Rigid Registration; Shape Correspondence; Anisotropy; Non-Isometry

Contents

Summary	iii
List of Abbreviations and Notations	x
List of Figures	xvii
1 Introduction	1
1.1 General topic	2
1.1.1 Goals	2
1.1.2 Motivation	2
1.1.3 Organisation	3
1.2 Contributions	4
1.3 Summary	5
2 Preliminary Concepts	6
2.1 Euclidean Space	6
2.2 Parametric Representation of Curves and Surfaces	7
2.3 Norms and Geodesics	8
2.4 Manifolds	10
2.5 Normals	13
2.6 Rigid registration	14
2.7 Neighbourhoods	15
2.8 Problems in practice	17
2.9 Summary	20
3 Literature Review	21
3.1 Scope	21
3.2 Previous surveys	22
3.3 Registration	23

3.3.1	Pre-processing	24
3.3.2	Correspondence computation	25
3.3.3	Transformation optimisation	26
	3.3.3.1 Data	26
	3.3.3.2 Regularisation	28
3.4	Shape correspondence	32
3.4.1	Local feature descriptors	32
3.4.2	Geodesic-based correspondence	32
3.4.3	Functional maps	33
3.4.4	Learning-based correspondence	35
3.4.5	Partial matching problem	37
3.4.6	Matching shape collections	38
3.5	Relevant approaches to this thesis	39
3.5.1	Traditional Non-Rigid ICP (N-ICP)	40
3.5.2	Non-Rigid Registration with Anisotropic Estimation	40
3.5.3	3D-CODED	40
3.5.4	Efficient Deformable Shape Correspondence via Kernel Matching	41
3.5.5	Reweighted Position and Transformation Sparsities	42
3.5.6	Genetic Isometric Shape Correspondence	42
3.5.7	RSDS Wrap 3.3 & 3.4	43
3.5.8	Partial Functional Correspondence	44
3.5.9	Deblurring and Denoising of Maps between Shapes	44
3.5.10	Continuous and Orientation-preserving Correspondences via Functional Maps	44
3.5.11	FARM+	45
3.5.12	CMH Connectivity Transfer	45
3.5.13	ZoomOut	46
3.6	Correspondence datasets	46
3.7	Non-isometry in datasets	46
	3.7.1 Dataset taxonomy	48
3.8	Evaluative measures	51
3.9	Medical image registration of skulls	53
3.10	Contributions	55
3.11	Summary	55

4	Non-rigid registration under anisotropic deformations	60
4.1	Introduction	60
4.2	Related work	63
4.3	Methodology	64
4.3.1	Algorithm overview	64
4.3.2	Anisotropy estimation	68
4.3.2.1	Local anisotropic deformation metric	69
4.3.2.2	Measuring anisotropic geodesic distance	70
4.3.3	Correspondence generation	71
4.3.3.1	Local feature matching	71
4.3.3.2	Correspondence identification based on automatic cor- respondence inference	72
4.3.4	Correspondence pruning	74
4.3.5	N-ICP with extended ARAP regularisation	75
4.4	Evaluation	76
4.4.1	Error measure	76
4.4.2	Implementation	77
4.4.3	Synthetic datasets	77
4.4.4	Real datasets	81
4.5	Limitations	82
4.6	Conclusions	82
4.7	Summary	83
5	SHREC’19: Shape correspondence with isometric and non-isometric deformations	88
5.1	Introduction	89
5.2	Dataset	90
5.2.1	Ground-truth construction	91
5.2.2	Evaluation	93
5.3	Correspondence methods	94
5.3.1	Traditional Non-Rigid ICP (N-ICP)	94
5.3.2	Non-Rigid Registration with Anisotropic Estimation	95
5.3.3	3D-CODED	96
5.3.4	Efficient Deformable Shape Correspondence via Kernel Matching	96
5.3.5	Reweighted Position and Transformation Sparsities	97
5.3.6	Genetic Isometric Shape Correspondence	97

5.3.7	RSDS Wrap 3.3	98
5.4	Results	98
5.5	Conclusions	101
5.6	Summary	103
6	SHREC'20: Shape correspondence with physically-based deformations	105
6.1	Introduction	105
6.2	Dataset	106
6.2.1	Initial correspondences	108
6.3	Correspondence methods	109
6.3.1	Robust Non-Rigid Registration with Reweighted Position and Transformation Sparsity	109
6.3.2	Partial Functional Correspondence	109
6.3.3	Non-rigid Registration under Anisotropic Deformations	109
6.3.4	Efficient Deformable Shape Correspondence via Kernel Matching	109
6.3.5	Deblurring and Denoising of Maps between Shapes	111
6.3.6	Continuous and Orientation-preserving Correspondences via Functional Maps	111
6.3.7	FARM+	111
6.3.8	R3DS Wrap 3.4	111
6.4	Results	112
6.4.1	Initial correspondences	112
6.4.2	Method comparison	113
6.5	Conclusions	115
6.6	Summary	115
7	SHREC'20: Shape correspondence with non-isometric deformations	124
7.1	Introduction	124
7.2	Dataset	127
7.2.1	Test sets	129
7.2.2	Initial correspondences	129
7.3	Correspondence methods	130
7.3.1	Baseline N-ICP	131
7.3.2	Non-rigid registration under anisotropic deformations	131
7.3.3	Robust Non-Rigid Registration with Reweighted Position and Transformation Sparsity	131

7.3.4	Efficient Deformable Shape Correspondence via Kernel Matching	131
7.3.5	Deblurring and Denoising of Maps between Shapes	131
7.3.6	Partial Functional Correspondence	131
7.3.7	Continuous and Orientation-preserving Correspondences via Functional Maps	132
7.3.8	CMH Connectivity Transfer	132
7.3.9	ZoomOut	132
7.3.10	R3DS Wrap 3.4	132
7.4	Evaluation	133
7.4.1	Error measure	133
7.4.2	Surface coverage measure	133
7.5	Results	136
7.5.1	Surface coverage	136
7.5.2	Geodesic error	138
7.6	Conclusion	144
7.7	Summary	145
8	Automatic orbital measures in CT scans using non-rigid registration	152
8.1	Introduction	152
8.2	Method	154
8.2.1	Data preparation (preprocessing)	155
8.2.2	Orbital landmark interpolation	156
8.2.3	Template selection	159
8.2.4	Surface registration	160
8.2.5	Quantitative measures	161
8.3	Results	162
8.3.1	Planarity	163
8.3.2	Asymmetry	165
8.4	Conclusions	166
8.5	Summary	167
9	Conclusions	168
9.1	Summary	168
9.2	Non-rigid registration under anisotropic deformations	169
9.3	SHREC'19: Shape correspondence with isometric and non-isometric deformations	169
9.4	SHREC'20: Shape correspondence with physically-based deformations	170

9.5	SHREC'20: Shape correspondence with non-isometric deformations	171
9.6	Automatic orbital measures in CT scans using non-rigid registration	171
9.7	Future work	172
9.7.1	Registration techniques	172
9.7.2	Datasets	173
9.8	Closing words	173
A	Citation count search terms	174
	Bibliography	175

List of Abbreviations and Notations

Abbreviations

ICP	Iterative closest point
N-ICP	Non-rigid ICP
SVD	Singular value decomposition
LBS	Linear blend skinning
MVC	Mean-value coordinates
GWAS	Genome-wide association study
ROC	Receiver operating characteristic curve
CMC	Cumulative match characteristic
CT	Computerised tomography
MRI	Magnetic resonance imaging
TPS	Thin-plate splines
FPFH	Fast Point Feature Histogram
SHOT	Signature of Histograms of Orientations
HKS	Heat Kernel Signature
WKS	Wave Kernel Signature
ARAP	As-rigid-as-possible
AUC	Area under the curve
FPS	Farthest point sampling

Notations

u	A vector
X, Y	Discrete shapes embedded in \mathbb{R}^3
\mathcal{M}, \mathcal{N}	Surfaces, 2-manifolds
$[0, \infty)$	Interval notation
$\mathcal{O}(\cdot)$	Big O notation
V, E, F	Vertices, edges & faces of X

List of Figures

1.1	An illustration of the hypothetical widget assembly machine.	1
2.1	An example of a parameterised differentiable curve.	7
2.2	An example of a differentiable surface.	8
2.3	An example of the physical representation of Euclidean distance with respect to a differentiable surface.	8
2.4	An example of various norms. The contour lines shown correspond to locations where the resulting value of l_p -norm is equal to one.	9
2.5	An example of a geodesic path on a differentiable surface.	10
2.6	An example of the tangent plane $T_p\mathcal{M}$ to a point p on the manifold \mathcal{M}	11
2.7	An example of isotropic deformation.	12
2.8	An example of anisotropic deformation.	12
2.9	Examples of simple low-genus shapes.	13
2.10	An example of a normal \mathbf{n} to a surface at point p	14
2.11	An example of a geodesic disc on a surface.	16
2.12	An example of (a) one-ring and (b) two-ring vertex neighbourhoods.	16
2.13	An example of (a) two surfaces (b) undergoing a topological change causing a thin handle. (c) The surface is cut at the shortest surface cycle to correct the topological error.	17
2.14	Common examples of non-manifold geometry.	17
2.15	An example of a cylinder intersecting a cube.	18
2.16	Regular polygons inscribed inside a circle to illustrate how a piecewise linear approximation of a circle improves as the number of edges increases [5]. The difference in area between the polygon and the actual circle's area is reported beneath each shape.	19

2.17	An example illustrating the difference between the shortest path following the edges of a mesh, and the shortest path over the faces. In (a) Dijkstra obtains the ℓ_1 -norm distance, while in (b) the correct distance is computed.	19
3.1	Number of published papers on the topic of non-rigid shape correspondence and registration. Based on Web of Science data, see Appendix A for further details about the search.	22
3.2	An illustration of the neighbourhood of (a) triangles, (b) spokes, and (c) spokes-and-rims.	28
3.3	Overview of the genetic algorithm [172].	43
4.1	An isotropic geodesic.	64
4.2	(a) Closest matching descriptor between X and Y , the groups of coloured squares represent feature vectors. (b) Pruning descriptors based on geodesic consistency, the boxes with integer values represent the distance to neighbouring points in X and the corresponding points in Y	65
4.3	N-ICP using a coarse set of hard correspondences. To the left, the source surface (X), represented by the red line, is non-rigidly aligned to the target surface (Y), in blue. This produces the deformed surface X' , in red, on the right.	66
4.5	An anisotropic geodesic.	66
4.4	The resulting pairs of eigenvectors/values of each corresponding point between X and Y based on the anisotropic non-isometric stretch.	67
4.6	Given the point x_2 on surface X , the geodesic compatibility of correspondences is measured for each candidate correspondence on Y (<i>i.e.</i> , y_2 & y_3). In this case, the distance values between y_2 and the known correspondences are more consistent with the distance values of x_2 , than the values of y_3	67

4.7	Illustrations that demonstrate the idea of estimating anisotropic non-isometric deformation. (a) An anisotropic deformation of points between two shapes (X & Y) may be approximately modelled by an ellipse. The surrounding unlabelled points represent neighbouring points on the source shape $\mathbf{s}_j \in N_X(\mathbf{s}_i)$ and the predicted correspondence of \mathbf{t}_v on surface Y . (b) Two pairs of eigenvectors ($\hat{\mathbf{e}}_{1a}, \hat{\mathbf{e}}_{2a}, \hat{\mathbf{e}}_{1b}, \hat{\mathbf{e}}_{2b} \in \mathbb{R}^2$) represent stretching for an edge \mathbf{v}_{ab} . (c) \mathbf{v}_{ab} transforming into new basis $\mathbf{B} = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$	71
4.8	Example problematic region between two shapes with no correspondences.	73
4.9	Exemplar illustration of geodesic distance. For a point at the centre of the leopard’s face, the geodesic distance to any given point on the surface is indicated by the local colour value.	77
4.10	TOSCA high-res dataset results.	78
4.11	Results for the TOSCA cat, gorilla, David & dog sets, (a) source/target models used, (b) combined registration error of our method, (c) combined registration error of [210].	79
4.12	Registration results on TOSCA dog. (a) initial pose source (top) target pose (bottom), (b) our method (c) Yang et al. [210].	79
4.13	Deformation results on Bouncing dataset [202]. (a) initial pose source, (d) target pose (bottom), (b,e) our method (c,f) sparse non-rigid registration method [210].	80
4.14	Results on SHREC’16 partial dataset compared with results (Partial Functional Maps & Random Forests) from the dense methods reported in [43]. The evaluation is split up into models containing (a) holes, and (b) cuts.	84
4.15	Results of our method in different configurations on SHREC’16 partial dataset. In configurations where PSF is off, standard geodesics are used, rather than anisotropic ones (skipping lines 7 & 8 of Algorithm 1). Where correspondence inference is off, we skip lines 10 & 11 of Algorithm 1. The evaluation is split up into models containing (a) holes, and (b) cuts.	85
4.16	Results of our proposed non-rigid registration technique on partial body scans [4]. (a,b,c) initial pose source X (blue), target Y (white), (d,e,f) overlapped result.	86
4.17	Registration results of pairs of models from the FAUST dataset.	86

4.18	Screenshots of correspondence results of (a,d) our method, (b,e) Li et al. [109] and (c,f) Vestner et al. [200] on pairs of inter-person shapes from the FAUST dataset (left source, right target). Colours represent correspondence between shapes (<i>i.e.</i> , the same point on each shape should have the same colour). Note the reflection of correspondences of [200] in (c), and the colour bleed at the intersection between legs in (f). The mean geodesic error for each scan pair: (a) 0.0373, (b) 0.0419, (c) 0.2971, (d) 0.0639, (e) 0.0820, and (f) 0.0840.	87
5.1	Examples of objects in the dataset.	88
5.2	Illustrations of some of the challenges in the dataset.	90
5.3	A photo of the real wooden hand used in the dataset after markers were drawn.	92
5.4	An illustration of the processing pipeline applied after extracting (a) a local patch is projected to a 2D plane, (b) a linear discriminator is applied to classify pixels, (c) morphological operations (<i>e.g.</i> , dilating and pruning) are then applied, (d) the final skeleton is then used to count the number of endpoints.	93
5.5	An illustration of the construction of a query string using the points neighbouring a query point (circled with two rings). The query string encodes the distance (by order), colour and number of spokes.	93
5.6	Shape pairs from test-sets 0, 2 & 3 with ground-truth correspondences visualised. Matching colours between shape pairs represent corresponding points.	94
5.7	This figure shows an example output of the deformed template using [74].	96
5.8	Extremely sparse (top) and sparse (bottom) correspondences produced by Sahillioğlu [172] on some pairs.	97
5.9	Results of all methods on each test set.	99
5.10	Results for all test sets.	103
5.11	Wooden hand object from the dataset coloured by the surface normals to illustrate the lack of high frequency geometric detail on the shape's surface.	104

6.1	Illustrations of the surfaces of meshes with different internal materials coloured by the surface normals. These materials help induce varying magnitudes and frequencies of protrusions and indentations on the surface.	107
6.2	Illustrations of the contents of the benchmark dataset. (a) partial (source) meshes, (b) the full/watertight (target) mesh. Note that the stretch pose with couscous filling was omitted in preliminary stages of scanning as its appearance was similar to the same pose with risotto filling. All scans were filled with 36 fl. oz. (or approx. 1065ml) of grains or beans, except for scans where the shape was inflated where they were filled with 56 fl. oz. (or approx. 1656ml).	117
6.3	A photo of the model (in the pose of scan00) used for the dataset with markers.	118
6.4	Examples of texture transfer using the ground-truth correspondences. (a) Target shape. (b-e) Source shapes. Correspondences were transferred and interpolated using a landmark-based correspondence method [67]. The original texture was projected onto the coronal (frontal) plane of the rabbit in (a).	118
6.5	Illustrations of some of the challenges in the dataset. (a) Partial scans (green indicates the boundary). (b) Complex deformations. (c) Missing geometry caused by self-occlusion.	118
6.6	Error curves for the methods used to establish an initial correspondence. Note, Pruning [192] only estimated $\approx 7.4250\%$ of correspondences.	119
6.7	Comparison of the error of methods initialised with different initial correspondence techniques (N-ICP [27], Pruning [192], SEG [100], or None). (a) R3DS Wrap 3. (b) Ezuz and Ben-Chen [67]. (c) Dyke et al. [55]. (d) Ren et al. [155].	120
6.8	Cumulative error curves with scans grouped by the pose exhibited. (a) twist (scan no. 1, 4 & 8). (b) indent (scan no. 2, 5 & 9). (c) inflate (scan no. 3, 6 & 10). (d) stretch (scan no. 7 & 11).	121
6.9	Cumulative error curves with scans grouped by the internal material selected. (a) Couscous (scan no. 1, 2 & 3). (b) Risotto (scan no. 4, 5, 6 & 7). (c) Chickpea (scan no. 8, 9, 10 & 11).	122
6.10	The performance of methods over all test sets.	123

7.1	A simple example of a homologous part between quadrupeds. The coloured area represents the matching part for the hind leg. Whilst inter-species, it is possible to intuitively recognise that these areas correspond.	126
7.2	Illustrations of the contents of the benchmark dataset simplified to consist of a maximum of 100,000 faces.	127
7.3	A screenshot of the software used by specialists to annotate corresponding points between shapes. The rhino on the left was initially labelled, and used as a reference for subsequent animals.	128
7.4	Illustrations of some of the challenges in the proposed dataset (a) partial scans (green indicates the boundary), (b) significant non-isometric deformations between pairs of models, and (c) topological inconsistencies: inherent to the original object or caused by scanning limitations.	128
7.5	An illustration of the initial correspondence computed by matching different descriptors after diffusion pruning.	130
7.6	Examples of a shape from the dataset segmented using increasingly dense sampling of Voronoi cells. Each sub-figure contains following number of Voronoi cells: (a) 2, (b) 10, (c) 25, (d) 50, (e) 150, & (f) 500.	135
7.7	An example illustrating a barycentric cell (the shaded area) for the vertex in the centre of a one-ring neighbourhood. The area of each vertex A_i —used in Alg. 2—also corresponds to the area of the vertex’s barycentric cell.	136
7.8	Examples of points distributed on the surface of a model from this dataset. (a) Bijective mapping, (b) part mapping (75.0%), (c) part mapping (50.0%), (d) sparse correspondence (50.0%), (e) sparse correspondence (25.0%), and (f) sparse correspondence (10.0%).	137
7.9	An example of the coverage measure computed using synthetic correspondences on a real mesh.	138
7.10	Overall coverage performance of methods that submitted results for all test sets.	138
7.11	Cumulative error curves for each test set, (a) test-set 0, (b) test-set 1, and (c) test-set 3.	139
7.11	Cumulative error curves for each test set, (d) test-set 3, and (e) test-set 4.	140
7.12	Overall performance of methods that submitted completed all test sets.	141
7.13	Cumulative error curves for each test set, (a) test-set 0, (b) test-set 1, (c) test-set 2.	142

7.13	Cumulative error curves for each test set, (d) test-set 3 and (e) test-set 4.	143
8.1	An illustrative diagram indicating the approximate location of the orbit. The dashed line represents the orbit on the human skull.	153
8.2	An overview of the proposed orbit registration pipeline.	155
8.3	An illustrative diagram highlighting the neck bone geometry that is typically excluded due to its separation from the skull, causing a disconnected mesh. The ruby area represents the neck supporting the human skull.	156
8.4	An illustration of the method used to measure the error between source point i and a discretised spline curve $(\dots, j - 1, j, j + 1, \dots)$	157
8.5	An illustration of a correctly projected discretised spline over a one-ring neighbourhood. All adjacent points typically share the same face or are on an adjacent face.	158
8.6	A registration error matrix illustrating the orbit alignment error of registering all manually labelled skulls against all of the other manually labelled skulls. Skull no. 3 & 5 achieve notably poor results both as the source and as the target scan. In the case of skull no. 3, the spline produced is misaligned. In the case of skull no. 5 the top of the skull is cut off just above the orbit, which is significantly inconsistent with the other skulls.	159
8.7	A comparison of the performance of using a barycentric interpolation method [91] vs. a naïve transportation method—nearest neighbour—on scans remeshed to 50,000 faces.	161
8.8	A histogram illustrating the planarity results of the predicted orbit registration pipeline using meshes with 50,000 and 100,000 faces. The mean of the ground-truth scans is also shown.	164
8.9	A histogram illustrating the asymmetry results of the predicted orbit registration pipeline using meshes with 50,000 and 100,000 faces. The mean of the ground-truth scans is also shown.	166

LIST OF FIGURES

Chapter 1

Introduction

Overview

What is registration? Generally, registration is the process of transforming two or more surfaces to embed them optimally in the same co-ordinate space. In the simplest case, we want to find a rotation and translation that will correctly align two surfaces in this shared space. Imagine an electronics assembly line in a factory that uses a robotic arm to place widgets on a *printed circuit board* (PCB). For the robotic arm to place the widget in the correct location, there must be an understanding of the location and orientation of the target PCB. By using a camera, or 3D scanner, placed above the conveyor belt, a registration algorithm positions a known template on top of the image, or scan, of the PCB. In this template fitting step, the template, or source, PCB has been registered to a target PCB.

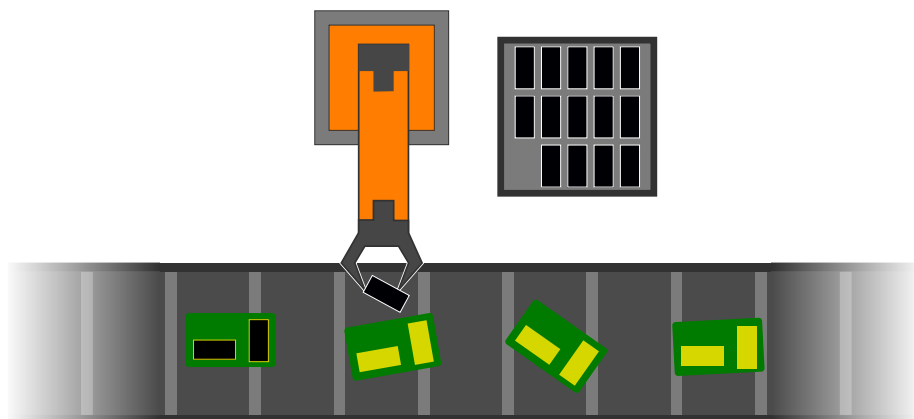


Figure 1.1: An illustration of the hypothetical widget assembly machine.

Interest in the area of geometry processing has been stimulated by the greater accessibility to hardware and software for capturing and reconstructing 3D shapes.

Tasks that require the comparison or combination of two or more shapes can benefit from accurate registration. The objective of this chapter is to provide a high-level overview of the work undertaken and contributions to the area of registration of non-rigidly deforming objects. The motivation and the structure of this thesis are described in Section 1.1. Section 1.2 then clarifies the key contributions of this thesis. Finally, a summary of the chapter is given in Section 1.3.

1.1 General topic

1.1.1 Goals

Broadly, the goal of this thesis is to develop, and contribute towards the further development of, registration and correspondence techniques for non-rigidly deforming 3D shapes. Specifically, this thesis focuses on the problem of isometry in the context of non-rigid registration. While aspects of surface registration have been studied thoroughly, a problem lacking research is that of non-isometric deformation; this is one of the primary focuses of this thesis. Furthermore, the tools and resources available for the comparison and measurement of the performance of existing approaches are insufficient. This has motivated the development of several datasets and benchmarking methodologies produced for this thesis.

In this thesis: a novel automatic registration pipeline is proposed; a set of resources for evaluating both registration and correspondence methods is developed; finally, a methodological contribution is made by utilising non-rigid registration on a large dataset to extract information, such as planarity and asymmetry, for automatic shape analysis.

1.1.2 Motivation

Non-rigid registration is a fundamental problem that underpins many applications in geometric processing (*e.g.*, template fitting [213], quadrangulation [14], surface reconstruction [223], classification [149], etc.). In the context of morphometry—the study of anatomical variations amongst objects—non-rigid registration enables the construction of statistical models, which may be used to conduct further study and develop an understanding of a group of objects.

The accessibility of high-quality scanning hardware continues to improve, with increasing amounts of 3D data produced every day. Accurately registering the data

manually is not feasible as—even for a small number of shapes—the process of registration is already infeasible. It is therefore crucial that automatic techniques that are sufficiently accurate are developed.

This thesis identifies that most current methods do not consider non-isometric deformations; therefore, a new technique is proposed. Additionally, a novel problem is examined where a set of labels must be transferred between shapes in a large database of 3D skulls. To address this problem, a modified non-rigid registration algorithm underlies a bespoke label transfer pipeline.

Furthermore, relevant evaluation procedures are necessary to determine how well registration methods perform in non-rigid scenarios, and enable comparisons between approaches. A distinct lack of suitable evaluative benchmarks for non-isometric deformations is identified. This leads to the development of three unique datasets.

1.1.3 Organisation

This thesis comprises of an introduction, literature review, five contribution chapters, and a conclusion.

The work undertaken and structure of this thesis is outlined in this chapter. Chapter 2 provides an introduction to a range of fundamental concepts that are necessary to understand the literature in this area.

The following chapter (Chapter 3) is a literature review. The chapter provides an overview of the current state-of-the-art in non-rigid registration, as well as discussing other relevant techniques, such as: functional mapping and deep learning correspondence methods. Additionally, an investigation into contemporary evaluation resources is conducted, producing a taxonomy of benchmark datasets for correspondence methods.

Chapter 4 introduces a novel non-rigid registration pipeline, which is designed to handle non-isometric deformations, achieving state-of-the-art performance.

Chapters 5, 6 & 7 propose novel correspondence benchmark datasets that consider various types of challenging deformations. These datasets were used in organised benchmarking challenges where multiple groups of researchers contributed correspondence results using either shape correspondence or registration methods. Chapter 5 presents a new dataset for evaluating shape correspondence. The benchmark is structured into a hierarchy, containing various types of challenging deformations (*e.g.*, piecewise-rigid, non-isometric, etc.). Chapter 6 is inspired by work on human perception. Four types of physically-based deformation are chosen, the coarseness of the internal material is varied, causing deformations for the same target pose to differ.

These present a challenge to both humans and automatic algorithms. In Chapter 7 specifically, it is observed that current evaluation measures are insufficient and only tell half the story about correspondence quality. Thus, a novel evaluation metric, which characterises the distribution and sparsity of correspondences is proposed.

Chapter 8 investigates the problem of automatically registering the orbits (*i.e.*, eye sockets) for a pair of skulls. A bespoke registration pipeline is developed. This method is applied to a large database of skulls to extract measures of orbital planarity and orbital asymmetry; thus, demonstrating the practicality of the approach.

Finally, this thesis is concluded in Chapter 9. Further discussions of the work undertaken and suggestions for future directions of work are identified.

1.2 Contributions

This thesis makes significant contributions to the problems of non-rigid registration and shape correspondence in the field of geometric processing. The key contributions of this work are summarised below.

Non-rigid registration

A novel non-rigid registration pipeline to handle large and anisotropic deformations is developed in Chapter 4. The contributions of the work are as follows:

- A novel method to estimate anisotropic deformations on a discrete mesh by using the principal scaling factor.
- A correspondence generation and correspondence pruning method based on local geodesics that copes with anisotropic deformations; this makes use of our anisotropic deformation estimate.
- The r -ring **ARAP** formulation is introduced for regularisation in non-rigid registration, which effectively handles challenging large deformations where existing registration methods fail.

In Chapter 8, a bespoke non-rigid registration pipeline is applied to a novel problem in which the asymmetry of the orbits of a collection of skulls is measured. The contributions of the work are as follows:

- A novel pipeline is proposed for the efficient localisation of orbits on large skull datasets.

- An evaluation of the accuracy of the proposed method compared against the performance of a sample of experts.

Benchmark datasets

Multiple datasets for the purposes of registration and correspondence evaluation are introduced.

The contributions of the work in Chapter 5 are as follows:

- Generation of a novel dataset of anthropomorphic shapes.
- A benchmark comprising of tiers based on the deformation exhibited, enabling greater insight into which scenarios a method may perform well in.

The contributions of the work in Chapter 6 are as follows:

- Generation of a novel dataset of physically-based deformations of a single object.
- Comprehensive use of the benchmark dataset to evaluate correspondence and registration methods in challenging full-to-partial scenarios.

The contributions of the work in Chapter 7 are as follows:

- Generation of a novel dataset of quadrupeds with sparse ground-truth correspondences labelled by three specialists.
- Development of a new measure to evaluate the coverage of correspondences on a shape’s surface—discussed further in Section 7.4.
- Systematic evaluation of the performance of a selection of recent shape correspondence methods, with additional quantitative insights into performance from our novel measure.

1.3 Summary

The objective of this chapter was to provide a high-level understanding of the work conducted in this thesis. Discussions about the goals and motivations provided some context regarding the work, this is expanded upon in each contribution chapter (Chapters 4, 5, 6, 7 & 8). To help guide the reader through this document, the precise organisation of the material was given. It is possible to skip through any of the contribution chapters—just listed—without losing too much context. The key contributions of this thesis were outlined here; however, the list is not exhaustive as minor contributions are omitted.

Chapter 2

Preliminary Concepts

Overview

The objective of the following section is to provide an introduction into the theoretical foundations of non-rigid registration for the novice reader, and also to ensure the information in this thesis is as self-contained as possible. In Section 2.1, Euclidean space, the space in which most of geometry problems are considered, is described. Section 2.2, provides a definition for the continuous representation of a curve. This definition is extended to describe a surface. Using this representation, Section 2.3 introduces the ℓ -norm and geodesics, which are commonly used to quantify distance extrinsically and intrinsically. Section 2.4 defines the Riemannian manifold, a geometric structure used to describe surfaces. Relevant topological properties (*i.e.*, genus) and geometrical properties (*e.g.*, non-isometry) are also described. An intuitive definition of the surface normal is given in Section 2.5. Section 2.6 discusses the problem of rigid registration; further discussion of registration methods is provided in Chapter 3. Section 2.7 introduces the concept of neighbourhoods. Section 2.8 discusses some common problems encountered in surface registration. Finally, Section 2.9 summarises the content of this chapter and makes recommendations on additional reading material.

2.1 Euclidean Space

Euclidean space is a real co-ordinate system that enables the position of a point to be determined. In geometry, conventionally, points are considered in the context of two- or three-dimensions. Two-dimensional spaces are typically represented by \mathbb{R}^2 where a point may be represented by $(x, y) \in \mathbb{R}^2$, while points in three-dimensional spaces are represented by $(x, y, z) \in \mathbb{R}^3$.

2.2 Parametric Representation of Curves and Surfaces

When approaching problems in the domain of geometric processing, the parameterisation of geometry is important to consider. This provides a mathematical grounding for a method; and practically, this approach ensures solutions generalise to problems regarding discrete meshes of varying density and non-uniformity.

Differentiable curves

Consider a point in a two-dimensional space (\mathbb{R}^2) at time unit u . The point's position at any time is given by the function $\mathbf{c}(u) = (\mathbf{x}(u), \mathbf{y}(u))$. As u changes, the point follows the path of a curve $\mathbf{c}(\cdot)$. More formally, a continuously differentiable curve may be expressed as follows [178]:

Definition 2.2.1 A parameterised differentiable curve in \mathbb{R}^n ($n = 2, 3, \dots$) is a continuous map $\mathbf{c} : I \rightarrow \mathbb{R}^n$, where $I \subset \mathbb{R}$ is an open interval (of end points $a, b \in [-\infty, \infty]$).




Figure 2.1: An example of a parameterised differentiable curve.

A differentiable curve is an example of a one-dimensional manifold. This is illustrated in Fig. 2.1.

Differentiable surfaces

Conceptually, a continuous surface embedded in \mathbb{R}^3 is not too dissimilar from a curve in \mathbb{R}^3 . A differentiable surface is an example of a two-dimensional manifold. Intuitively, while a curve comprises of a single parameter (u), a surface is comprised of two parameters. *i.e.*, :

Definition 2.2.2 Let S be a surface embedded in \mathbb{R}^3 . A parameterised differentiable surface in \mathbb{R}^3 is a continuous map $\Phi : S \rightarrow \mathbb{R}^3$, where $S \subset \mathbb{R}^2$.

This concept is illustrated in Fig. 2.2.




Figure 2.2: An example of a differentiable surface.

Point $p = (u, v)$ may be mapped from the local co-ordinate system \mathbb{R}^2 of surface S to the global co-ordinate system \mathbb{R}^3 by the co-ordinate map:

$$\Phi(u, v) = (\mathbf{x}(u, v), \mathbf{y}(u, v), \mathbf{z}(u, v)).$$

2.3 Norms and Geodesics

The *norm* is a measure of the length of a vector. The most familiar norm is the ℓ_2 -norm, which measures Euclidean distance, see Fig. 2.3.




Figure 2.3: An example of the physical representation of Euclidean distance with respect to a differentiable surface.

Euclidean distance is the length of a line between two points in a Euclidean space. Euclidean distance is an extrinsic measure. This is measured by the *Euclidean norm*

or the ℓ_2 -norm, the square root of the sum of all squared elements, *i.e.*, the Euclidean norm of an n -dimensional vector $\mathbf{u} = (u_1, \dots, u_n)$ is:

$$\|\mathbf{u}\|_2 = \sqrt{u_1^2 + \dots + u_n^2}$$

There are several other norms¹ that are commonly used in computational geometry, these are listed below:

Name	Definition
ℓ_0 -norm	The number of non-zero elements in \mathbf{u} <i>i.e.</i> , $\sum_{i=1}^n u_i^0$, assuming $0^0 = 0$.
ℓ_1 -norm	The sum of the absolute values of its elements <i>i.e.</i> , $\sum_{i=1}^n u_i $.
ℓ_2 -norm	The squared sum of values of its elements <i>i.e.</i> , $(\sum_{i=1}^n [u_i^2])^{\frac{1}{2}}$.
ℓ_∞ -norm	The maximum of the absolute values of its elements <i>i.e.</i> , $\max_{i=1}^n u_i $.

The contour plot in Fig. 2.4 illustrates how the response of ℓ_p -norm changes as p is varied in \mathbb{R}^2 . An interactive example of this is available².




Figure 2.4: An example of various norms. The contour lines shown correspond to locations where the resulting value of ℓ_p -norm is equal to one.

A *geodesic*, or a *geodesic path*, is the shortest path between two points on the surface of a shape, while *geodesic distance* is a measure of length of a geodesic path. Geodesic distance is an intrinsic measure. This is illustrated in Fig. 2.5.

¹This concept generalises to the form called ℓ_p -norm where $\|\mathbf{u}\|_p = (\sum_{i=1}^n [u_i^p])^{\frac{1}{p}}$.

²Interactive plot: <https://www.desmos.com/calculator/h7wtnao1x4>




Figure 2.5: An example of a geodesic path on a differentiable surface.

More formally, a path on a surface is considered to be a *geodesic* when:

Definition 2.3.1 Let $\mathbf{c} : I \rightarrow S$ be a parameterised differentiable curve, where $I \subset \mathbb{R}$. If the line joining any two points on a curve $\mathbf{c}(a)$ and $\mathbf{c}(b)$ is the shortest path of all curves between these points, then this curve is called a geodesic.

2.4 Manifolds

A *manifold* is any geometric structure that has a locally Euclidean topology. A *smooth manifold* is any manifold that is infinitely differentiable.

A curve (Definition 2.2.1) is classed as a one-dimensional manifold, and a surface (Definition 2.2.2) is classed as a two-dimensional manifold.

Riemannian manifolds

A *Riemannian manifold* (\mathcal{M}, g) comprises of a smooth differentiable manifold \mathcal{M} and a Riemannian metric g .

Riemannian metric

The *Riemannian metric* is an important tool used in studying differential geometry. It is used to measure angles and lengths of curves on smooth manifolds that are classified as *Riemannian manifolds*.

Definition 2.4.1 Given a manifold \mathcal{M} , $T_p\mathcal{M}$ is the tangent space of \mathcal{M} at point p . The Riemannian metric g comprises of the set of inner products (given by the dot product in \mathbb{R}^2) of $T_p\mathcal{M}$.

This concept is illustrated in Fig. 2.6.




Figure 2.6: An example of the tangent plane $T_p\mathcal{M}$ to a point p on the manifold \mathcal{M} .

Geodesics on Riemannian manifolds

Based on Definition 2.4.1 of a Riemannian metric, a geodesic distance may be measured as $g(\mathbf{c}(a), \mathbf{c}(b))$. To simplify the notation, unless stated otherwise, from hereon we shall exclude \mathbf{c} and assume the shortest path is always followed.

Isometry

If two Riemannian manifolds \mathcal{M} and \mathcal{N} have a smooth mapping $\Phi : \mathcal{M} \rightarrow \mathcal{N}$, such that the geodesic distance between points is preserved, then Φ is an *isometric mapping*.

More formally, given a mapping between two Riemannian manifolds, *i.e.*, $\Phi : (\mathcal{M}, g_{\mathcal{M}}) \rightarrow (\mathcal{N}, g_{\mathcal{N}})$, Φ is considered to be an isometric mapping, if for any pair of points $(a, b) \in \mathcal{M}$ where $(\Phi(a), \Phi(b)) \in \mathcal{N}$, then

$$g_{\mathcal{M}}(a, b) = g_{\mathcal{N}}(\Phi(a), \Phi(b)).$$

Also, a subset between two manifolds $S_{\mathcal{M}} \subset \mathcal{M}$ and $S_{\mathcal{N}} \subset \mathcal{N}$ may be considered a *locally isometric mapping*, if for any pair of points $(a, b) \in S_{\mathcal{M}}$ where $(\Phi(a), \Phi(b)) \in S_{\mathcal{N}}$, then

$$g_{S_{\mathcal{M}}}(a, b) = g_{S_{\mathcal{N}}}(\Phi(a), \Phi(b)).$$

If Φ is locally isometric and diffeomorphic³, then Φ is *isometric*.

³Both Φ and Φ^{-1} are smooth and differentiable

Near-isometry

A mapping $\Phi : (\mathcal{M}, g_{\mathcal{M}}) \rightarrow (\mathcal{N}, g_{\mathcal{N}})$ may be considered to be ε -isometric if the pairwise distance between points $(a, b) \in \mathcal{M}$ map to $(\Phi(a), \Phi(b)) \in \mathcal{N}$ is

$$|g_{\mathcal{M}}(a, b) - g_{\mathcal{N}}(\Phi(a), \Phi(b))| \leq \varepsilon.$$

In practice, mappings tend not to be exactly isometric. ε -isometric mappings where ε is small are considered to be *near-isometric*. Selecting an appropriately small ε is application dependent, it may be hard to define and is typically not specified in most works.

Isotropy and anisotropy

When discussing deformation, *isotropy* refers to the quality of a given manifold to stretch uniformly, see Fig. 2.7. Local isotropic deformation occurs when a surface locally stretches in all directions by the same amount. A balloon being inflated is an intuitive example of isotropic expansion. While similarly, a balloon deflating is an example of isotropic contraction.




Figure 2.7: An example of isotropic deformation.

In contrast, *anisotropy* refers to deformations that non-uniformly stretch a manifold in different directions, see Fig. 2.8. A real-world example of anisotropic deformation can be observed by pulling on an elastic band. The stretch may be described by two orthogonal eigenvectors/values λ_1 and λ_2 . An isometric deformation will induce no stretching, *i.e.*, $|\lambda_1| = 1$ and $|\lambda_2| = 1$.




Figure 2.8: An example of anisotropic deformation.

Genus

Genus may be considered to be the number of topological holes in a surface. The genus of the shape is the same as the number of handles, see Fig. 2.9. Topological holes differ from geometric holes, which are caused by missing geometry. In cases where there are no geometric holes, the geometry may be referred to as *watertight*.




Figure 2.9: Examples of simple low-genus shapes.

The genus of a surface is closely related to the Euler-Poincaré characteristic. The theory considers a graph-based representation of a manifold, where $|V|$, $|E|$, $|F|$ are the numbers of vertices, edges, and faces respectively, the genus g can be calculated by rearranging the following equation:

$$\chi(g) = |V| - |E| + |F| = 2 - 2g.$$

In the case of shapes consisting of multiple connected components, the formula may be extended to incorporate $|B|$ number of boundary edges, and $|C|$ number of connected components, to compute g^* the total genus for all components:

$$|V| - |E| + |F| = 2|C| - 2g^* - |B|.$$

2.5 Normals

Given a locally diffeomorphic (*i.e.*, locally flat) surface. The normal may be described by a unit vector \mathbf{n} that is perpendicular to the surface at point p . Given p and any two arbitrary points a and b on the tangent plane—assuming the points are non-collinear—the normal may be computed using the cross product. This is illustrated in Fig. 2.10.

$$\mathbf{n} = \frac{(a - p) \times (b - p)}{\|(a - p) \times (b - p)\|_2} \quad (\text{where } p \neq a, p \neq b, a \neq b).$$




Figure 2.10: An example of a normal \mathbf{n} to a surface at point p .

2.6 Rigid registration

The general objective of rigid registration methods is to find a transformation that optimally aligns two sets of points. The rigid transformation may be described by a rotation R and translation t between a source shape \mathcal{M} and target shape \mathcal{N} .

Given a set of corresponding points between two shapes $(x_i, y_i) \in \mathcal{M} \times \mathcal{N}$, the following least squares equation is minimised:

$$E(R, t) = \sum_{i=1}^n \|(Rx_i + t) - y_i\|_2^2.$$

The resulting transformation may be applied to all points in the source shape to transform the correspondences to align with the target shape. $y_i = Rx_i + t$.

Rigid Iterative Closest Point registration

Iterative closest point (ICP) is an algorithm used to estimate the transformation between two shapes with unknown correspondence. Typically, a set of candidate correspondences that may be initially inaccurate is estimated. As the name suggests, an iterative approach aims to improve the accuracy of these correspondences. Two main steps are alternated between:

1. For each point on the source shape, find the closest point on the target shape.
2. Estimate a transformation between the set of corresponding points.

In subsequent iterations, after the first, the estimated transformation is applied to the source shape; this is used to find a new set of correspondences. Traditional **ICP** implementations assume surfaces are coarsely aligned. Depending on the complexity

of the geometry being aligned, distant initial orientations of surfaces cause **ICP** to fail.

The performance of each iteration is significantly improved through the use of k -D trees [18]. Given a source shape \mathcal{M} with n points and a target shape \mathcal{N} with m points. The cost of computing the closest points between the shapes in each iteration is $O(n \log_2 m)$.

The performance of each iteration of **ICP** is dependent on the strategies used to compute a correspondence between shapes.

Common strategies to align points use *point-to-point* [19] and *point-to-plane* [39] error metrics.

$$E_{\text{point-to-point}} = \sum_{i=1}^n \|(Rx_i + t) - y_i\|_2^2.$$

$$E_{\text{point-to-plane}} = \sum_{i=1}^n (((Rx_i + t) - y_i) \cdot \mathbf{n}_i)^2.$$

These may be combined to help estimate transformation using a parameter (*e.g.*, α) to weight each component, *i.e.*,

$$E = E_{\text{point-to-point}} + \alpha E_{\text{point-to-plane}}.$$

Low [123] produced a technical report that goes into further details of the implementation of point-to-plane-based **ICP**.

Further variants of **ICP** proposed in literature seek to improve the stability to noise and outliers, accuracy, and speed of convergence [90, 165].

2.7 Neighbourhoods

When performing analyses on geometry, information about neighbourhoods is often required to compute elementary properties, such as surface normals and curvature.

Geodesic r -neighbourhoods

In topology, a neighbourhood is typically thought of as a disk or ball that has a particular radius $r \in \mathbb{R}^+$. Anything encompassed by the ball at a particular point (a) is considered to be within the neighbourhood of a . More formally:

$$B(a, r) = \{b \in \mathcal{M} \mid 0 < g_{\mathcal{M}(a,b)} < r\}.$$

This is illustrated in Fig 2.11


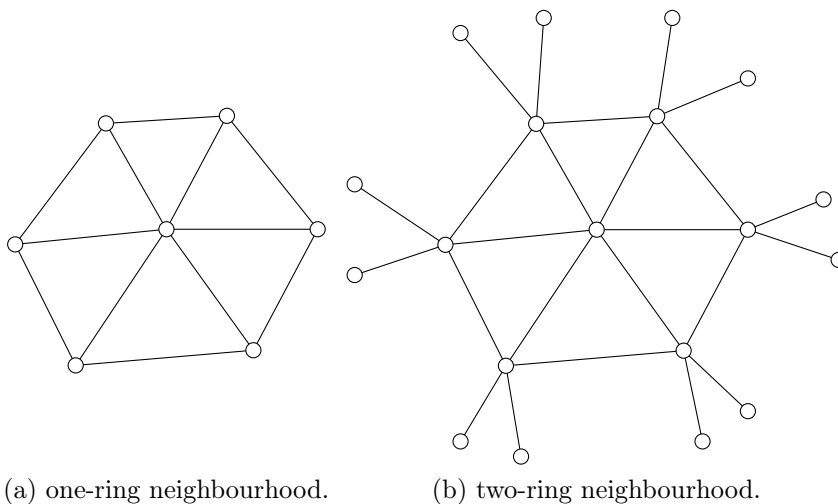


Figure 2.11: An example of a geodesic disc on a surface.

This concept is compatible with discrete meshes, however computing geodesics can be expensive in terms of time. This cost also increases as r increases.

r -ring graph neighbourhoods

On discrete meshes, the surface may be treated as a graph. Neighbourhoods are commonly defined by vertices that are connected to a particular vertex (a) within a certain number of edge hops r (note that here $r \in \mathbb{N}_+$). This method is efficient to compute for small discrete neighbourhoods. This is illustrated in Fig. 2.12.



(a) one-ring neighbourhood.

(b) two-ring neighbourhood.

Figure 2.12: An example of (a) one-ring and (b) two-ring vertex neighbourhoods.

Neighbourhoods may also be described directly by edges, or faces. The weakness of this approach comes from the underlying assumptions that the shape has a good quality triangulation and that each vertex is reasonably spaced. This is often not the case when using real data, meaning that additional mesh refinement is required.

2.8 Problems in practice

Capturing data

Depending on the scanning technology used, different limitations may be observed. All scanning methods have a degree of error. These present in the form of *geometric noise* and *topological change*.

Handles that are incorrectly created during reconstruction are due to parts of a surface being incorrectly fused or punctured. These can be corrected by either cutting (or pinching) thin handles, or collapsing (or filling) small holes [204]. Such strategies are essential for correspondence and registration methods that require input shapes to have a low or consistent genus. This concept is illustrated in Fig. 2.13.




Figure 2.13: An example of (a) two surfaces (b) undergoing a topological change causing a thin handle. (c) The surface is cut at the shortest surface cycle to correct the topological error.

Geometry quality

Non-manifold geometry Non-manifold geometry refers to geometry that is not manifold everywhere. Such geometry is often undesirable in real-world capture scenarios and is problematic for many tasks. Properties such as volume cannot be easily measured on non-manifolds. This problem has been extensively investigated [130]. Fig. 2.14 illustrates two common non-manifold scenarios.




Figure 2.14: Common examples of non-manifold geometry.

Self-intersection The intersection of geometry occurs when a part of two or more shapes occupy the same space, such as in Fig. 2.15. In the real world, such anomalies are not possible. However, errors during the scanning and reconstruction of an object can cause these to occur. Geometry that is free from self-intersections is important for applications such as 3D printing. Self-intersections can be rectified using the method proposed by Zhu et al. [221]. For non-rigid registration algorithms, typically the deformation phase is under-constrained, accepting solutions with self-intersections. In pose estimation, an additional collision term is used to address self-intersection [16]. Recently, demonstrated by Eisenberger et al. [62], it is possible to efficiently compute a deformation field that does not introduce additional self-intersections for shape correspondence.




Figure 2.15: An example of a cylinder intersecting a cube.

Discretisation error

As a shape’s representation coarsens, the fidelity decreases, as exemplified in Fig 2.16. Processing geometry with large numbers of polygons may be infeasible for some algorithms even on modern hardware. Therefore, obtaining an accurate—yet low-polygon—discretisation of a shape is important.

For example, coarsening strategies have been used to compute smaller wave kernel maps for functional mapping methods, which often require large amounts of memory [105, 120].

Geodesics

Many geometry processing methods require the computation of distance. The complexity of the renowned graph-based shortest path method *Dijkstra’s algorithm* [52] is $\mathcal{O}(|E| + |V| \log |V|)$ in the worst case. For graphs, this method is considered highly efficient and is appropriate for many graph-based problems. However, for meshes, this approach tends to lead to a severe overestimation of the actual distance, see




Figure 2.16: Regular polygons inscribed inside a circle to illustrate how a piece-wise linear approximation of a circle improves as the number of edges increases [5]. The difference in area between the polygon and the actual circle’s area is reported beneath each shape.

Fig. 2.17. This is because the true geodesic often passes through the mesh faces, rather than precisely following the edges of the mesh. Computing exact geodesics that cut through faces, rather than following edges, can be expensive; therefore approximate methods are often employed. For the computation of exact geodesics the prominent method—self-titled—“continuous Dijkstra” [135] is used with similar performance to [52]. This method is precise but slow on fine meshes. For large-scale applications approximate approaches are favoured. The *Fast Marching method* [97] exploits the eikonal equation to estimate the shortest path across faces. The original method performs in $\mathcal{O}(|V| \log |V|)$.




Figure 2.17: An example illustrating the difference between the shortest path following the edges of a mesh, and the shortest path over the faces. In (a) Dijkstra obtains the ℓ_1 -norm distance, while in (b) the correct distance is computed.

2.9 Summary

The aim of this chapter has been to prepare the uninitiated reader for the topic of surface registration. Although this chapter's contents cannot be comprehensive, it was to provide an adequate basis for the discussions in the subsequent chapters of this thesis. Suggested further reading include "Polygon Mesh Processing" [26] for an overview of a variety of geometry processing techniques, and "Numerical Geometry of Non-rigid Shapes" [32] for a deep exploration of non-rigid shape analysis. Going forth, the rest of this thesis discusses work that contributes novel findings to the area of non-rigid registration.

Chapter 3

Literature Review

Overview

This chapter presents a review of notable works that are relevant to the technical contributions of this thesis. The review aims to provide an overview of the current state-of-the-art, identifying general trends and notable deficiencies of these techniques—the scope of which is defined in Section 3.1. Section 3.2 introduces works that have surveyed relevant literature in the areas of surface registration and shape correspondence. In Section 3.3, the general surface registration pipeline is decomposed into its key components (*i.e.*, pre-processing, initial correspondence, transformation optimisation). Notable surface registration techniques that appear in the literature are then discussed in the relevant subsection. An overview of a wide variety of shape correspondence approaches is given in Section 3.4. Section 3.5 describes in detail relevant registration and shape correspondence methods that are used later in evaluative work. In Sections 3.6 & 3.7, an overview and discussion of suitable datasets for the evaluation of non-rigid registration and shape correspondence methods is given. Section 3.8 surveys evaluative techniques used to measure the performance of registration methods. Section 3.9 discusses a task-specific set of registration methods that are designed to register skulls with non-rigid deformation. The overall contributions of this chapter are discussed in Section 3.10. Finally, the content of the chapter is summarised in Section 3.11.

3.1 Scope

Principally, this thesis considers the problem of non-rigid registration. The non-rigid registration is deconstructed into its constituent parts, and notable extensions or modifications proposed in the literature are discussed. Due to the broad applicability of

the dataset contributions of this thesis, this chapter also provides an introduction to other relevant problems and emerging areas (*e.g.*, functional mapping, learning-based correspondence, and correspondence benchmark datasets). While not necessarily addressing the registration problem directly, these approaches establish meaningful correspondences between non-rigid shapes, which may be subsequently used to initialise a registration method. Only techniques that are designed to handle either point cloud or 3D triangulated mesh structures are included. A related problem that is outside the scope of this review is reconstruction. Many reconstruction techniques rely on registration to accurately produce a 3D model from a series of partial scans. These methods often adopt simple yet efficient registration techniques that are capable of being incorporated into online reconstruction pipelines.

The field of shape correspondence has an abundance of literature. A brief exploration (Fig. 3.1) shows that in the past decade there has been a marked increase of publications related to this problem.




Figure 3.1: Number of published papers on the topic of non-rigid shape correspondence and registration. Based on Web of Science data, see Appendix A for further details about the search.

3.2 Previous surveys

Surveys provide a broad introduction to newcomers and are helpful to experts as they may offer a new perspective on a given problem.

A seminal survey in the field of shape correspondence by van Kaick et al. [196] (2011), provides a broad introduction and organisation of numerous solutions. Additional

discussions of evaluation methods and applications for shape correspondence are presented. Tam et al. [191] (2013) investigate a broad range of techniques for both rigid and non-rigid registration. The survey deconstructs the principal components of notable registration pipelines. Further distinction between works is made by the constraints incorporated to simplify the complexity of the registration problem. Interestingly, despite the age of the surveys, van Kaick et al. [196] and Tam et al. [191] still remain relevant with many of the concepts discussed in these works still prevalent in state-of-the-art methods used to date.

More recently, Ovsjanikov et al. [146] (2017) provide a thorough overview of the functional map concept, demonstrating the applicability of the functional map representation for relevant problems. At each stage of the functional map pipeline, the work introduces and discusses various solutions that seek to address a specific challenge. Sahillioğlu [171] (2019) provides a much needed comprehensive overview of many recent shape correspondence methods, complementing [196]. This survey extends the general shape analysis taxonomy by Biasotti et al. [20] (2016), categorising recent correspondence methods by additional relevant criteria.

For shape descriptors, Rostami et al. [163] (2019) distinguish between shape descriptors by their application by constructing a taxonomy. Shape correspondence and shape similarity are reviewed by Bronstein et al. [34] (2017), providing further context in the field of deep learning for two particular applications. For shape similarity a descriptor is only necessary to discriminate between shapes globally; whereas, for shape correspondence, the ability of a descriptor to localise matching parts of a shape is essential.

3.3 Registration

This section discusses existing literature regarding the registration of non-rigidly deforming surfaces.

Point set registration A related problem is the registration of non-rigid *point sets*—rather than meshes. Zhu et al. [220] survey the many current works on point set registration. Often such methods adapt either **ICP** [19, 39] or use a probabilistic approach with a mixture model [138]. While relevant, these methods do not consider valuable geometric data contained in mesh-based shapes; therefore, they do not seek to solve the same problem. This thesis focuses greatly on mesh-based registration

where surfaces have been discretised into triangles. Since some of the contributions discussed in this thesis are underpinned by **ICP**, we shall discuss the approach here.

From rigid registration The problem of aligning two sets of data points has been studied for a long time in the statistical domain. In the last three decades, there has been an increasing interest in the alignment of different co-ordinate spaces. Besl and McKay [19] and Chen and Medioni [39] introduced **ICP**, a simple iterative approach to registration for point sets. The method is computationally efficient. Many variants of the original **ICP** formulation have since been proposed. The general **ICP** framework is as follows: First, a binary (hard) correspondence is established using a closest point criterion, such as nearest neighbours. Then a rigid transformation between the two surfaces is estimated using a least-squares formulation. These two steps are repeated until a convergence criterion is met. The method has also been extended to handle non-rigid cases. The main shortcoming of this approach is the necessity of a good initial alignment. When handling large-scale deformations, typical (non-rigid) **ICP** often fails. Further discussion and evaluation of various implementations of **ICP** have been conducted by Rusinkiewicz and Levoy [165].

Non-rigid registration Non-rigid approaches were initially developed to register range scans where the scanning devices used had been incorrectly calibrated [108]. Typical non-rigid registration frameworks may be deconstructed into a few distinct components that all contribute toward solving the registration problem:

- Pre-processing;
- Correspondence computation;
- Transformation optimisation.

3.3.1 Pre-processing

Many efficient registration methods that are capable of processing large numbers of points in a short amount of time have been proposed [37, 109, 109, 88].

For computationally intensive approaches, where processing is impractically slow for large meshes, coarser representations are required. Often, either farthest point sampling strategies [92] or mesh simplification [70] are used (see [200, 55]). Recent work by Liu et al. [120] seeks to preserve spectral properties rather than the geometric properties for tasks such as shape matching. Conversely, for some approaches,

additional subdivision of particularly coarse meshes may be required. In the case of [66], the technique is underpinned by the embedding method by Aigerman and Lipman [3]. However, the method [3] fails to produce a valid embedding if two or more geodesic paths share an edge. This can happen at the boundary of a mesh when the triangulation is too coarse.

In the case of [200], input is further cleaned using [12] to construct a watertight representation.

Methods that rely on specific topologies (see [171]) or limited geometric error require meshes to be cleaned in advance. Although not directly addressed by modern registration pipelines, in the case of many methods, non-manifold geometry, erroneously high genera, and disconnected geometry must all be fixed beforehand. Automatic approaches have been proposed, however there are no guarantees that such fixes will preserve important information.

For many shape matching and registration methods some properties that are invariant, or are treated as invariant, when finding a correspondence can be pre-computed. Typical registration methods [19, 27] do not manipulate the target shape and therefore do not require properties (*e.g.*, surface normals) of the target model to be re-computed at each optimisation step. For both source and target shapes, often, feature descriptors and geodesics can be computed just once. Feature descriptors are used to measure properties of a mesh that are invariant to certain deformations; therefore, re-computing descriptors on the deformed mesh has little benefit, and is likely to introduce greater errors. In the case of [55], an estimation of anisotropy is updated upon each iteration, causing geodesics to be updated to reflect the estimated anisotropic distortion. Commonly, k -d tree structures are employed to efficiently estimate the closest points between a query point and a point set. The initial data structure may be constructed once for the target shape, then repeatedly queried with points from either the source or deformed shape.

3.3.2 Correspondence computation

Rusinkiewicz and Levoy [165] discuss a variety of procedures to compute a set of correspondences between rigid shapes. The authors discuss strategies to select key-points, finding correspondences between keypoints, weighting pairs, and strategies to reject pairs. These approaches also apply to non-rigid registration methods.

Initial correspondences [71, 82, 192] compute an initial set of sparse correspondences. First, feature descriptors for each point are computed. These are often based on purely geometric data; as [113] demonstrates, when a shape’s texture is available it is possible to incorporate this additional information. Next, a subset of correspondences that preserve their global coherence are matched through an optimisation procedure. These methods rely heavily on the robustness of extrinsic or intrinsic distance measures. When used to initialise a registration method, these approaches tend to improve results [210, 109, 55]. However, results may worsen if particular deformations (*e.g.*, non-isometries) are not considered by the model, or if the shape descriptors used are sensitive to such changes. Recently, Kleiman and Ovsjanikov [100] proposed a method to co-segment pairs of shapes that exhibit non-isometric deformation. The corresponding regions were used to initialise a functional mapping method [144].

Correspondences determined during optimisation The majority of methods find the closest point from a source to a target shape. This can be done efficiently by constructing a k -d tree [165]. For registration problems where a sub-vertex accuracy is necessary, source points may be projected onto the surface of the target scan. The projected points are then used as targets for fitting—this approach is used in Chapter 8. Papazov and Burschka [147] discovered that for registration, applications simply using nearest neighbours can lead to many undesirable matches. Papazov and Burschka [147] and Jiang et al. [87] find the closest point on the target surface that is suitably close and has a compatible normal with a similar orientation. Then, among the neighbours of this point, the point that has the closest distance to the average distance of the neighbours is selected.

3.3.3 Transformation optimisation

When estimating the optimal transformation, the problem is treated as an energy optimisation in which a data and a regularisation term—or set of terms—is minimised (*i.e.*, $E = E_{\text{data}} + E_{\text{reg}}$). This subsection discusses various notable formulations for both terms.

3.3.3.1 Data In the context of rigid registration, a variety of popular error measurement strategies are discussed by Rusinkiewicz and Levoy [165]. Also, Babin et al. [15] compare the performance of different functions for penalising fitting errors. The authors conclude that the ℓ_1 -norm obtains reasonably good results, with

no additional parameters required to tune. Here, literature surrounding two popular strategies (point-to-point and point-to-plane) is discussed.

Point-to-point Commonly, formulations for point-to-point methods use an ℓ_2 -norm-based energy term [6, 82, 27, 208, 210, 55]. This is suitable in some cases, but has been shown to be less suitable in cases where there are significant amounts of noise and outliers.

$$E_{\text{point-to-point}} = \sum_{(i,j) \in C} \|\mathbf{x}'_i - \mathbf{y}_j\|_2^2,$$

where $(i, j) \in C$ specifies the closest point \mathbf{y}_j on Y from \mathbf{x}'_i .

There are several notable variations of this formula. Li et al. [107] also use ℓ_2 -norm, but limit the search space to principally the target shape’s surface, with the energy term modified to approximately align regions with no overlap. This approach provides good results for scans with small holes or missing parts. Alternatively, Li et al. [109] and Jiang et al. [88] use ℓ_1 -norm. This approach is considered to help alleviate some undesirable overfitting caused by ℓ_2 -norm measures. To make the error term more representative on non-uniform meshes, Pauly et al. [148] use an ℓ_2 -norm-based measure with the distance weighted by the area of each vertex. Yao et al. [212] observe that, while less sensitive to noise, alternative formulations that incorporate the ℓ_p -norm ($p < 1$) [28] are costly to compute, requiring the use of slower minimisation techniques. Instead, Yao et al. [212] formulate the error metric using Welsch’s function [80], which can be minimised efficiently.

Point-to-plane The seminal work by Chen and Medioni [39] proposed the use of a point-to-plane algorithm for registration. In cases of surfaces with sparse high-frequency details, this approach enables surfaces to slide along each other, obtaining a superior convergence rate compared to early point-to-point methods [165]. Since then, the point-to-plane measure has continued to be used in more recent works (*e.g.*, [123, 82, 27, 208, 55]).

$$E_{\text{point-to-plane}} = \sum_{(i,j) \in C} \|(\mathbf{x}'_i - \mathbf{y}_j) \cdot \mathbf{n}_i\|_2^2,$$

where \mathbf{n}_i is the respective normal of \mathbf{x}_i .

Zhang et al. [217] incorporate Welsch’s function [80] for solving rigid registration problems. However, as well as utilising a different optimisation procedure, here it is formulated as a point-to-plane term.

3.3.3.2 Regularisation Non-rigid deformations can have a high number of degrees of freedom. It is therefore important to consider how registration may be constrained. Depending on the scenario, undesirable deformations can be penalised. Here we discuss some notable regularisation formulations that may be considered when developing a registration method.

Cell edge sets: triangles, spokes, spokes-and-rims A cell refers to the local neighbourhood N , cells are used to compute local surface properties. Spokes refer to edges (j, k) connected to a subject vertex x_i , this is denoted as $(j, k) \in N_i$. Rims refer to edges connected to the vertices that neighbour a subject vertex. Fig. 3.2 illustrates the vertices and edges that each type of neighbourhood comprises of.




Figure 3.2: An illustration of the neighbourhood of (a) triangles, (b) spokes, and (c) spokes-and-rims.

Triangles: The simplest expressible neighbourhood is the three edges of a triangle f_i . This neighbourhood information of discrete faces F is explicitly defined in most common data representations as,

$$\hat{N}_i = \{(x_i, x_j), (x_j, x_k), (x_k, x_i) \mid (x_i, x_j, x_k) \in F\}.$$

Spokes: The neighbourhood of a given vertex $x_i \in V$ is the set of edges E that are adjacent with x_i .

$$N_i = \{(x_j, x_k) \in E \mid x_i = x_j\}.$$

To increase the size of the neighbourhood from a one-ring neighbourhood to a two-ring neighbourhood (N_i^2), this formulation is extended to include the edges of neighbours from the first ring:

$$N_i^2 = \bigcup_{j \in N_i} N_j.$$

Spokes-and-rims: The neighbourhood of a given vertex $x_i \in V$ is the set of edges that are adjacent with x_i combined with the edges between vertices that share an edge with x_i .

$$\tilde{N}_i = \{(x_{ja}, x_{ka}), (x_{jb}, x_{kb}), (x_{jc}, x_{kc}) \in E \mid \\ x_i = x_{ja} \text{ and } x_i = x_{jb} \text{ and } x_{ka} \neq x_{kb} \text{ and } x_{ka} = x_{jc} \text{ and } x_{kb} = x_{kc}\}.$$

Isometric or length-preserving Geodesic distances on the surface should remain unchanged. Such preservation can also be incorporated by combining angle-preserving and scale-preserving constraints. Kilian et al. [94] ensure isometric shape deformation by preserving geodesics between surface points. Sauvage et al. [177] propose a deformation framework for B-spline surfaces that incorporates a volume constraint.

Given a one-ring neighbourhood ($r = 1$), Sorkine and Alexa [185] seek an optimal rotation \mathbf{R}_i between a set of points \mathbf{x} and a corresponding deformed set \mathbf{x}' . Chen et al. [38] generalise this approach to r -ring neighbourhoods

$$E_{\text{ARAP}} = \sum_{(j,k) \in N_i^r} w_{jk} \left\| (\mathbf{x}'_k - \mathbf{x}'_j) - \mathbf{R}_i(\mathbf{x}_k - \mathbf{x}_j) \right\|_2^2,$$

where w_{jk} is some edge weighting scheme (*e.g.*, co-tangent weights). The one-ring formulation of this term is commonly used for non-rigid registration. Dyke et al. [55] incorporate an r -ring term into a registration pipeline to provide stronger regularisation to cope with large-scale deformation.

Levi and Gotsman [106] incorporate an energy term that penalises the difference between neighbouring rotations. This leads to smooth bending.

$$E_{\text{SR-ARAP}} = E_{\text{ARAP}} + \alpha E_{\text{SR}},$$

where

$$E_{\text{SR}} = \sum_{(j,k) \in N_i} w_{jk} \left\| \mathbf{R}_k - \mathbf{R}_j \right\|_2^2.$$

Sumner et al. [188] describe deformation with affine transformations—here, $\hat{\mathbf{R}}$ represents an affine matrix. To penalise undesirable shearing a term that promotes the affine matrix to appear as a rotation matrix is used.

$$E_{\text{rot}} = \sum_{i=1}^{|V|} \text{Rot}(\hat{\mathbf{R}}_i),$$

where

$$\text{Rot}(\hat{\mathbf{R}}) = (\mathbf{a}_1 \cdot \mathbf{a}_2)^2 + (\mathbf{a}_1 \cdot \mathbf{a}_3)^2 + (\mathbf{a}_2 \cdot \mathbf{a}_3)^2 + (1 - \|\mathbf{a}_1\|)^2 + (1 - \|\mathbf{a}_2\|)^2 + (1 - \|\mathbf{a}_3\|)^2.$$

In this case, $\hat{\mathbf{R}}$ can be an affine transformation, and \mathbf{a}_1 , \mathbf{a}_2 , and \mathbf{a}_3 are 3×1 column vectors of $\hat{\mathbf{R}}$. This approach for registration is used by Li et al. [107]. Brown and Rusinkiewicz [35] also propose an efficient non-rigid registration method based on thin-plate spline deformation, which may be separated into affine and non-affine parts.

Amberg et al. [6] use a similar approach to E_{SR} for an affine matrix $\hat{\mathbf{R}}$. Adopting a weighting matrix $\mathbf{G} = \text{diag}(1, 1, 1, \gamma)$, the energy term is applied to every edge $(i, j) \in E$:

$$E_S = \sum_{(j,k) \in E_i} w_{jk} \left\| (\hat{\mathbf{R}}_j - \hat{\mathbf{R}}_i) \mathbf{G} \right\|_2^2.$$

This formulation allows the amount of shearing to be controlled by γ . In the implementation described by Amberg et al. [6], this regularisation term is iteratively relaxed. It should be noted that the over-relaxation of this term may lead to overfitting. The problem of determining the ideal weight of the regularisation term or terms is universal for all registration methods.

Conformal, similar or angle-preserving These formulations permit isotropic deformation, but penalise angular distortion.

Yamazaki et al. [208] use an *as-similar-as-possible* (ASAP) energy term in a non-rigid ICP pipeline to register range scans.

$$E_{\text{ASAP}} = \sum_{(j,k) \in N_i^r} w_{jk} \left\| (\mathbf{x}'_k - \mathbf{x}'_j) - s_i \mathbf{R}_i (\mathbf{x}_k - \mathbf{x}_j) \right\|_2^2,$$

$s_i \in \mathbb{R}^+$ is a scale factor.

Inspired by Levi and Gotsman [106], Jiang et al. [87] incorporate the E_{SR} term

$$E_{\text{CASAP}} = E_{\text{ASAP}} + \alpha E_{\text{SR}}.$$

This additional term constrains deformations to bend smoothly, but also permits isotropy.

Yoshiyasu et al. [214] propose to solve a linear approximation of the *as-conformal-as-possible* (ACAP) formulation. Both ASAP and ACAP induce conformal deformations, the key difference between the two formulations is in the optimisation calculation. Comparisons using a linear approximation of ACAP find that this leads to

varied results, in terms of quality and convergence [214].

$$E_{\text{ACAP}} = \sum_{i=1}^{|V|} \text{Conformal}(\mathbf{R}_i),$$

where

$$\begin{aligned} \text{Conformal}(\mathbf{R}) &= (\mathbf{a}_1 \cdot \mathbf{a}_2)^2 + (\mathbf{a}_1 \cdot \mathbf{a}_3)^2 + (\mathbf{a}_2 \cdot \mathbf{a}_3)^2 \\ &+ (\|\mathbf{a}_1\| - \|\mathbf{a}_2\|)^2 + (\|\mathbf{a}_2\| - \|\mathbf{a}_3\|)^2 + (\|\mathbf{a}_3\| - \|\mathbf{a}_1\|)^2. \end{aligned} \quad (3.1)$$

Wu et al. [206] incorporate E_{ARAP} and E_{ACAP} and E_{SR} terms into a registration pipeline. The combination of these regularisation terms help constrain deformation during optimisation to establish accurate correspondences between isometric shapes.

Deformation constrained by Laplacians was proposed by Sorkine et al. [184]. This approach penalises heterogeneous transformations of vertices with respect to the centre of mass of their cell. This leads to deformations being smooth across the surface.

$$E_{\text{Laplacian}} = \sum_{i=1}^{|V|} \|\mathbf{T}_i(X')\delta_i - L(\mathbf{x}'_i)\|_2^2,$$

where the uniform Laplacian operator is formulated as

$$L(\mathbf{x}_i) = \frac{1}{|N_i|} \sum_{(j,k) \in N_i} (\mathbf{x}_j - \mathbf{x}_i).$$

Equiareal or scale-preserving Equiareal maps penalise distortion in the area of a deformation.

$$E_{\text{Equiareal}} = \sum_{i=1}^{|V|} \|\text{Area}(\mathbf{x}'_i) - \text{Area}(\mathbf{x}_i)\|_2^2,$$

where $\text{Area}(\cdot)$ is the area of a given cell. For a given point, where $\sum_{(i,j,k) \in N_i^r}$ is a triangle, $\text{Area}(\cdot)$ may be calculated as follows:

$$\text{Area}(\mathbf{x}_i) = \frac{1}{6} \sum_{(i,j,k) \in N_i^r} |(\mathbf{x}_j - \mathbf{x}_i) \times (\mathbf{x}_k - \mathbf{x}_i)|.$$

Jiang [86] tetrahedralise the interior of shapes in order to form a volume-preserving mapping. By applying either E_{ARAP} or E_{ASAP} to both cells on the internal points, as well as the mesh surface, it is possible to constrain the deformation to penalise changes in volume. However, tetrahedralisation becomes problematic on open meshes

(*e.g.*, partial scans, scans with holes, etc.) where the volume boundary is not defined. For suitably complete meshes, Barill et al. [17] propose an approximate form of the winding number to determine the interior and exterior of a mesh.

Popa et al. [154] propose a deformation model capable of handling anisotropic materials through the use of local co-ordinate frames. Achenbach et al. [1] successfully apply an anisotropic regularisation energy with edge Laplacians to a template face model, to improve alignment for anisotropic facial regions such as wrinkles.

Our main observation is that recent novel registration techniques have limited their focus to optimisation of these constraints, with less focus on other aspects of the pipeline—*e.g.*, correspondence selection. With respect to the work undertaken in this thesis, a novel approach that unifies diffusion pruning and registration is developed to address non-isometric deformation.

3.4 Shape correspondence

Shape correspondence methods attempt to compute the correspondence between two shapes without applying transformations to align the source and target meshes [196].

3.4.1 Local feature descriptors

Local feature descriptors may be used to establish correspondences between keypoints with similar local geometry. The majority of shape correspondence methods rely on matching descriptors that are invariant under specific deformations. A non-extensive survey of notable approaches is given here, see the survey by Rostami et al. [163] for a more complete overview.

Over the years, many robust feature descriptors have been proposed [89, 69, 168, 169, 194]. They have been extensively evaluated by Guo et al. [76]. Among those, the *Signature of Histograms of Orientations* (**SHOT**) [194] is a highly popular descriptor, because of its computational efficiency and state-of-the-art performance. More recently, Zeng et al. [216] proposed a volumetric feature descriptor using a *Convolutional Neural Network* (CNN) architecture. Other descriptors based on CNNs have also been proposed [203, 136]. However, these techniques may still generate inconsistent correspondences and usually require a large amount of data for training.

3.4.2 Geodesic-based correspondence

Geodesics are often used to measure the distortion between candidate correspondences [196, 193, 82, 192, 55, 9].

Huang et al. [82] present a technique capable of registering large-scale deformations. Near-rigidly deforming regions of a shape are segmented into clusters, which are then deformed and registered. The method thus assumes a mostly piecewise-rigid deformation from the source to the target shape. To obtain reliable correspondences, they adapt the spectral matching technique by Leordeanu and Hebert [104], which computes a globally consistent set of correspondences by assuming a near-isometric deformation.

Vestner et al. [201] construct a probabilistic framework to estimate the location of a correspondence. Geodesic Gaussian kernels smooth existing sparse, or noisy, correspondence data. The authors find that the problem can be treated as a linear assignment problem, and an auction algorithm is employed. The approach enables the method to handle non-isometric deformations; however, these meshes must be bijective. Vestner et al. [200] improves upon [201], introducing the use of heat kernels and support for partial correspondence.

Sahillioğlu and Yemez [173] measure isometric distortion on a set of points on the source and target shapes using geodesics. Due to the computational complexity of this approach on shapes with a large number vertices, a sparse set of points is used. Sahillioğlu [172] extends this method using a genetic algorithm-based approach to find an optimal correspondence configuration of the points, and a new sampling technique. The sparse set is updated each iteration using an adaptive sampling technique to help further decrease the distortion error.

Arbel et al. [9] use geodesics to extract pairs of patches from a set of initial correspondences. The similarity of each pair of patches is measured by the uniqueness of feature matches within a patch, and the consistency of the distances between correspondences and a centre point. Correspondences are refined using a coarse-to-fine approach, where the localisation of a candidate correspondence is refined to ensure maximum similarity in a patch at a given scale. The use of Fast Point Feature Histogram (**FPFH**) [169] as feature descriptors makes the method robust to intrinsic symmetries; however, in cases where significant non-isometric or topological change is present, this method may fail.

3.4.3 Functional maps

The seminal work by Ovsjanikov et al. [144] proposes a shape matching technique that is computed in the functional domain. The main advantage of this method is the efficiency of solving the correspondence problem in a truncated functional map

representation. Functional maps have demonstrated great flexibility, with the facility to incorporate either point-to-point or region-based correspondences [100, 155].

The original method had several major deficiencies. [144] was incapable of distinguishing between intrinsically symmetric solutions. This meant that for shapes with a single intrinsic line of symmetry, there was a 50% chance that the correct solution would be chosen. The original approach used isometrically invariant descriptors, these may become unreliable under non-isometric deformation. Coupled with a pointwise recovery method that was only suitable for isometries, the method was limited to only (near-)isometric deformations. Furthermore, functional maps produce what is called a soft map, many reliant applications require hard—point-to-point—correspondences. The original work employed a simple **ICP**-based method to recover a point-to-point representation. However, this approach was limited to isometric shape pairs and could not specify correspondences with a sub-vertex precision.

Symmetries In certain classes of shape there tends to be a significant amount of intrinsic symmetry. This is often a problem for principally intrinsic approaches, which fail to discriminate symmetric solutions. For benchmarks [32, 43, 159], symmetric ground-truths are available. For methods that are not robust to symmetric matches, two symmetric solutions may be submitted for each shape pair, and the solution with the lowest error is selected. This has led to subsequent works that incorporate information about the surface to address this. As well as two input shapes, [145] requires a known symmetry map as input.

For orientable surfaces, it is possible to automatically disambiguate symmetries with the use of appropriate shape descriptors (*e.g.*, **FPFH** [169]). This has been demonstrated in the context of intrinsic shape correspondence [110]—albeit not a functional mapping approach. Recent work has incorporated information about surface normals directly into the optimisation formula [155]. Yoshiyasu et al. [215] identify symmetries by constructing low-dimensional embeddings of input shapes.

Non-isometry The original method [144] was proposed to address isometric deformation, however a plethora of follow-up works have demonstrated possible extensions that help facilitate non-isometric deformation [95, 162, 101, 65, 67, 155, 156, 181, 66, 133, 85, 63]. A subset of these methods are assessed in this thesis.

Pointwise map recovery In the original work [144], a strategy similar to **ICP** was employed to recover point-to-point correspondences in the spectral domain through

the use of nearest neighbours. Ezuz and Ben-Chen [67] note that this approach is not suitable in non-isometric scenarios. Rodolà et al. [162] present a recovery method that incorporates a Coherent Point Drift method [138] to align the spectral embeddings of the two shapes. However, the method was only capable of coping with bijection (*i.e.*, both shapes have the same number of points, all of which correspond uniquely). A follow-up work by the same authors [161] extended this method to handle partial matching problems. However, this approach was limited and could not specify correspondences with a sub-vertex precision. Ezuz and Ben-Chen [67] and Ezuz et al. [68] propose point-to-face recovery methods capable of handling non-isometric shapes. The methods apply a series of optimisation techniques to ensure maps remain smooth and reversible. Both approaches are highly flexible, being applicable to most functional mapping pipelines, at handling non-isometric deformation problems.

3.4.4 Learning-based correspondence

Lately, there has been an increasing interest in utilising deep learning on meshes, with the growing development of deep learning tools that may be applied intrinsically to the mesh surface [34]. As discussed below, a large proportion of literature in this area aims to develop robust feature descriptors, which are designed to handle domain/task/class-specific deformations.

Bronstein et al. [34] give an overview of deep learning methods. Each technique is broadly classified by the domain the network works in (*i.e.*, frequency, spatial, or spatial-frequency domain). To avoid repetition of the existing survey, this section provides a broader overview of learning-based correspondence techniques.

Shallow learning Wuhrer et al. [207] and Salazar et al. [175] propose correspondence methods to establish correspondences between human models and face scans respectively. The pipeline used by both works is similar. An initial set of labelled models are used to train a Markov random field to predict landmark locations on an input scan. Landmarks between the template and input scan can then be coarsely aligned. In the case of the human models, a template fitting method [78] is then used to determine the human’s pose, while in the case of the face scans, a linear blendshape model is used to estimate the appearance of a face. Finally, the deformed template and input scan are aligned using non-rigid registration. This approach is adequate for specific well-constrained problems. However, outlier poses that are not encoded in the original shape model may be problematic.

Generating training data For facial correspondence, Gilani et al. [72] use a deep neural network architecture to extract predefined landmarks on the input scan. The network is trained using a set of synthetically generated models—with intrinsic ground-truths—produced by a commercial solution. Many other notable learning-based works [74] take a similar approach, using existing parameterised models (*e.g.*, [8, 121, 153, 224]) to generate models.

Metric learning Sun et al. [190] note that in the case of non-isometric deformations, local metrics are dissimilar, causing typical spectral descriptors in the Euclidean space to fail. Sun et al. [190] seek to learn an optimal embedding for spectral descriptors in a new metric space, which promotes invariance for task-specific deformations.

Feature learning Many deep learning methods seek to develop accurate pointwise descriptors that are invariant to a specific class of objects. By training a network using weaker descriptors, such as Heat Kernel Signature (**HKS**) [189] or Wave Kernel Signature (**WKS**) [13], it is possible to learn an improved descriptor with greater specificity [118]. In early works, functional mapping methods (*e.g.*, [144, 158]) were initialised by these newly learnt feature descriptors [158, 129, 25]. More recently, state-of-the-art refinement techniques (*e.g.*, [201, 133]) have been proposed that are initialised by new learning-based descriptor methods [136, 199, 36, 77, 53]. These methods consistently obtain improved results on domain-specific problems when compared with state-of-the-art methods that use hand-engineered feature descriptors.

Litany et al. [116] develop a learning-based method that focused specifically on the problem of shape correspondence. Unlike previous works, the method incorporates [144] directly into the network’s training scheme for calculating loss. Roufosse et al. [164] and Halimi et al. [77] extend this method to facilitate unsupervised learning—without the use of ground-truths—by optimising constraints such as bijectivity, isometry, and orthogonality during training. Unsupervised methods that seek to preserve isometric mappings fail to be robust against non-isometries and topological changes [77].

Wei et al. [203] propose a learning-based correspondence method that is robust to areas with no overlap, occlusions and partial matches. The method learns feature weights using an end-to-end pipeline on depth maps of human models from three datasets.

Data scarcity Classical deep learning approaches typically perform well when a large amount of quality training data is available. As evidenced in Section 3.6, there are a limited number of shape correspondence datasets that exist—with many limited to just anthropocentric shapes. Alternatively, when valuable labelled data is scarce, it is possible to utilise either a Siamese [24, 116, 164, 77, 75, 53, 190] or triplet [36] neural network for training the network to produce consistent features. Siamese networks are often trained by penalising descriptors [81] or maps [116, 53] that produce matches in non-corresponding regions. This is done through the formulation of a loss term that measures embedding feature space distance [81], metric/map distortions [116] or other intrinsic properties of the functional map [53].

The development of effective deep learning techniques on shape correspondence problems is still a maturing area. Current methods demonstrate exceptional results for class-specific deformations; however, these methods do not adequately generalise to other deformation scenarios. There are further avenues of research yet to be explored to address acquisition (*e.g.*, occlusion, topological change, etc.), as well as practical implementation challenges (*e.g.*, resource efficiency). An ensemble-based approach combining multiple existing models with the use of a multi-branch network might lead to fast convergence of optimal feature weights in unsupervised scenarios.

3.4.5 Partial matching problem

Partial meshes often occur in real-world scenarios, where parts of the object being scanned are occluded from the scanning equipment, and thereby are missing from the final reconstruction. Several real-world datasets with ground-truths that contain partial meshes have been proposed [157, 21, 57, 58]. The amount of overlap between the source and target scans varies between these datasets.

Rodolà et al. [159, 160] present a variant of the synthetic dataset by Bronstein et al. [32] that contains partial meshes. Shapes were modified to include holes and cuts that introduce different types of challenges. The dataset provides insights into the performance of shape correspondence or registration techniques on varying degrees of partial shapes (between 20% and 80%). Lähner et al. [102] modify Rodolà et al. [158] to produce a dataset with geometric fusions that reflect the limitations of common real-world scanning techniques.

Many methods have been developed that are capable of addressing the problem of establishing correspondences between a partial scan and a watertight model [152, 170, 139, 115, 159, 200, 54, 9, 55, 205]. However, these approaches are not designed to handle scans that have little overlap.

Li et al. [107] address this with a method that identifies holes and occlusions in the target scan. An energy fitting term is subsequently re-weighted to avoid incidentally penalising the source from aligning to such regions.

A shape descriptor designed to handle partial matches for both correspondence and retrieval tasks is proposed by van Kaick et al. [197]. The authors demonstrate good performance under topological change, despite the method’s reliance on geodesics. However, the method is still sensitive to changes that have a more extreme effect on geodesic measures.

To address partial scans, intrinsic approaches impose assumptions about the possible types of deformation; typically, this excludes non-isometry [152, 170, 139, 115, 159, 54, 9, 205]. For many real-world scenarios this limitation is not adequate. There is a clear need for methods that are designed to handle partial matching problems for non-isometrically deforming shapes.

3.4.6 Matching shape collections

Traditionally, the problem of establishing correspondences has been posed as a pairwise problem, in which just two shapes are analysed. However, we are often interested in establishing correspondences amongst a collection of shapes rather than just two. By collectively optimising correspondences, partially incorrect mappings can be identified and rectified using the rest of the set. It is therefore possible to establish accurate and consistent maps amongst a set of shapes. For tasks such as co-segmentation, this is necessary.

Initial work by Davies et al. [50] sought to construct 3D statistical shape models while also establishing a homeomorphic set of correspondences between a collection of shapes automatically. The motivation behind the proposed method was to address the time-consuming process of manually labelling correspondences to build shape models. Similar to the same authors’ research into 2D statistical shape models [49], the problem of establishing correspondences is treated as a minimisation of the *description length* used for a spherical parameterisation of each shape. The proposed objective function seeks to reduce the complexity (*i.e.*, description length) of the parameterisations, which is manipulated by diffeomorphic transformations. The method is limited to genus 0 shapes. In experimental results, the method was shown to be able to produce a dense consistent parameterisation based on a sparse set of automatic landmarks.

Since this thesis focuses on mappings between pairs of shapes, notable facets of existing literature for this problem are only briefly discussed.

Nguyen et al. [141] note that when encountering intrinsically symmetric shapes, it is often non-trivial for pairwise correspondence methods to resolve these cases. Typically, matching methods for collections of shapes incorporate a cycle-consistency criterion to refine weak maps in the collection. The authors describe a method to produce cycle-consistent maps between shapes. The consistency and accuracy are measured using geodesic distances between corresponding points, which is globally optimised. The main limitation of this approach is that it assumes that only a few incorrect maps exist in the collection.

Huang et al. [83] and Kim et al. [96] use fuzzy approaches to solve the problem. Huang et al. [83] present a method that is capable of refining a sparse set of initial pointwise correspondences. The method uses a two-step iterative optimisation procedure in which soft correspondences are produced and refined before computing a set of point-to-point correspondences, while Kim et al. [96] uses a diffusion-based approach to establish a set of fuzzy correspondences between a collection of shapes in an iterative manner. Huang et al. [84] use a functional map-based approach to establish correspondences amongst heterogeneous shapes.

Cosmo et al. [45] introduce a new method that does not require initial pointwise correspondences. Furthermore, the method is capable of establishing consistent correspondences between partially similar shapes, disregarding regions that do not match. However, Cosmo et al. [45] relies on stable geodesics, and the method may discard good matches between corresponding regions that are only present in a minority of shape pairs in favour of globally consistent matches.

Recently, Cohen and Ben-Chen [41] propose a method that relies on pairwise matching to establish consistent correspondences between two collections of shapes.

One challenge that occurs when using pairwise methods to establish correspondences over a collection of shapes is selecting a representative model. As discussed in Chapter 8, it is possible to find a suitable model by measuring the correspondence error of a sub-sample of shapes.

3.5 Relevant approaches to this thesis

Several existing surface registration and shape correspondence methods are examined in the benchmark dataset Chapters 5, 6 & 7. Each method is described in moderate detail here. The exact parameters used in the subsequent experiments are detailed later on in the respective chapters.

Classification Based on the comprehensive survey paper by Sahillioğlu [171], all methods evaluated in Chapters 5, 6 & 7 have been categorised in Table 3.1 based on the criteria described in [171]. Please refer to the original survey for the precise definition of each criterion.

3.5.1 Traditional Non-Rigid ICP (N-ICP)

Bouaziz and Pauly [27] propose an implementation of **N-ICP** that extends the original rigid formulation of ICP [19] to support non-rigid deformation. The method repeatedly applies the following two steps until convergence. In the first step, it finds correspondences between surfaces based on closest point matching, similar to ICP. In the second step, point-to-point distances are combined with point-to-plane distances to minimise the distance between correspondences. The combination of distance measures speeds up convergence, aligning the surfaces. It also uses one-ring as-rigid-as-possible regularisation to smooth local deformations.

3.5.2 Non-Rigid Registration with Anisotropic Estimation

The method by Dyke et al. [55], which is described in greater detail in Chapter 4, follows the **N-ICP** framework that alternately improves correspondences and local transformations. The initial correspondences are obtained based on matching of local geometric features (**SHOT** [194] is used). In order to address (local) anisotropic deformations, the method iteratively estimates local anisotropy—represented as local principal directions and principal scaling factors. These estimates are then incorporated in an extended diffusion pruning framework [192] to identify consistent correspondences, taking anisotropy into account when calculating geodesic distances. Local regions with substantial stretching may end up with very few correspondences identified due to changes of local geometric features. To cope with significantly different input shapes, the method introduces additional correspondences by incorporating established correspondences as landmarks.

3.5.3 3D-CODED

The method by Groueix et al. [74] takes a deep learning approach for matching deformable shapes, and introduces Shape Deformation Networks which jointly encode 3D shapes and correspondences. This is achieved by factoring the surface representation into (i) a template that parameterises the surface, and (ii) a learnt global feature vector that parameterises the transformation of the template into the input surface.

By predicting this feature for a new shape, correspondences between this shape and the template are implicitly predicted. These correspondences can be improved by an additional step which improves the shape feature by minimising the Chamfer distance between the input and transformed template.

To learn a transformation between shapes, an encoder-decoder architecture is trained end-to-end to optimise a regularised reconstruction loss. 3D shape correspondences between two shapes X and Y are found by first using the decoder to compute the parameters that deform the template to each of the two shapes. For each point on shape X , its nearest neighbour is found on the template deformed to X . This template point has a known corresponding point in the template deformed to Y , which is then used to find the nearest neighbour in Y .

3.5.4 Efficient Deformable Shape Correspondence via Kernel Matching

Kernel Matching applies the method proposed by Vestner et al. [200]. The algorithm solves a series of linear assignment problems (LAPs) of the form

$$\mathbf{P}^{(k+1)} = \arg \max_{\mathbf{P} \in \Pi_n} \langle \mathbf{P}, \alpha \mathbf{F}_Y \mathbf{F}_X^\top + \mathbf{K}_Y^{t_k} \mathbf{P}^{(k)} \mathbf{K}_X^{t_k} \rangle, \quad (3.2)$$

$$\mathbf{P}^{(0)} = \arg \max_{\mathbf{P} \in \Pi_n} \langle \mathbf{P}, \alpha \mathbf{F}_Y \mathbf{F}_X^\top \rangle, \quad (3.3)$$

where \mathbf{P} is a permutation matrix, $\mathbf{F}_X, \mathbf{F}_Y$ are matrices of pointwise descriptors and $\mathbf{K}_X^{t_k}, \mathbf{K}_Y^{t_k}$ are the positive-definite heat kernel matrices with diffusion parameter t_k on shapes X and Y , respectively. Intuitively, the first term in Eqn. 3.2 describes descriptor similarity and the second how well the neighbourhood information is preserved by comparing heat kernels. By decreasing the diffusion parameter more emphasis is put on local neighbourhoods (cf. the ε - δ definition of continuity). The paper gives more details about connections to quadratic assignment problems of the form

$$\arg \max_{\mathbf{P} \in \Pi_n} E(\mathbf{P}) = \arg \max_{\mathbf{P} \in \Pi_n} \langle \mathbf{P}, \alpha \mathbf{F}_Y \mathbf{F}_X^\top + \mathbf{K}_Y^{t_k} \mathbf{P} \mathbf{K}_X^{t_k} \rangle, \quad (3.4)$$

as well as interpretations in terms of *kernel density estimation* and *low pass filtering of correspondences*.

By construction the algorithm yields bijections and is—in its basic variant—only applicable to pairs of shapes with consistent sampling, in particular with the same number of vertices. In addition solving the LAP becomes intractable for high number

of vertices. Thus a multi-scale approach is applied that overcomes both issues: At each scale a subset of the vertices is sub-sampled (with increasing density). The solution of each scale induces Voronoi cells in the following scale, and sparse initial correspondences for the next scale.

Since the Voronoi cells can directly be put into correspondence, the vertices are matched between corresponding cells and thus a set of smaller LAPs is solved instead of one big LAP. Notice that the payoff matrices (Eqn. 3.2) of an LAP capture the correspondences between the centres of all Voronoi cells. The different number of vertices within corresponding cells can be tackled via slack variables. As a consequence the final matching is a bijection between subsets of the vertices on X and Y . The fraction of unmatched vertices is low and tends to appear in areas with inconsistent sampling.

3.5.5 Reweighted Position and Transformation Sparsities

In order to cope with challenges of non-rigid registration, namely high degrees of freedom and presence of noise and outliers, Li et al. [109] propose a robust non-rigid registration method using reweighted sparsities on position and transformation to estimate the deformations between 3D shapes. Observing that large position and transformation errors tend to concentrate on local areas, which can be considered as sparse signals over surfaces, they formulate the energy function with position and transformation sparsity on both the data term and the smoothness term, and define the smoothness constraint using local rigidity. The double sparsity based non-rigid registration model is enhanced with a reweighting scheme to further improve its robustness. The formulation is solved by transferring it into four alternately-optimised sub-problems which have exact solutions and guaranteed convergence. To cope with large differences in source and target shapes, diffusion pruning [192] is applied to obtain initial correspondences based on matching of local **SHOT** features [194], and further correspondences are introduced during iterative optimisation based on closest points, similar to the standard N-ICP framework.

3.5.6 Genetic Isometric Shape Correspondence

The method by Sahillioğlu [172] exploits the permutation creation ability of genetic optimisation to find the permutation matrix that encodes correspondences between two point sets. To this end, Sahillioğlu [172] provides a genetic algorithm for the 3D shape correspondence problem. The point sets to be matched are sampled from

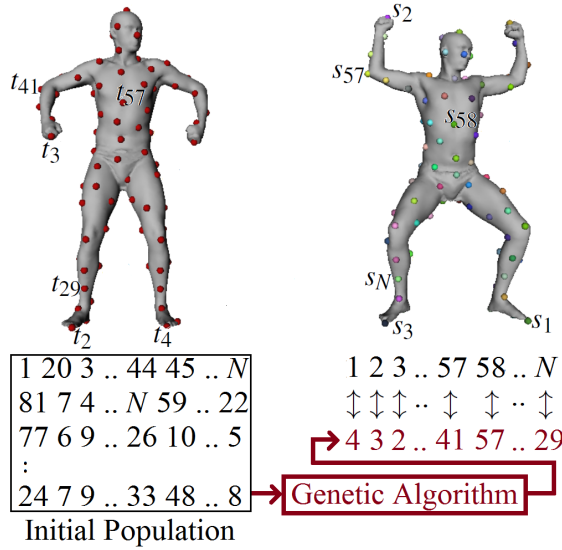


Figure 3.3: Overview of the genetic algorithm [172].

two isometric (or near-isometric) shapes. The sparse one-to-one correspondences produced by this algorithm minimise the following isometric distortion function:

$$\mathcal{D}_{\text{iso}}(\phi) = \frac{1}{|\phi|} \sum_{(x_i, y_j) \in \phi} \frac{1}{|\phi'|} \sum_{(x_i, y_m) \in \phi'} |d_g(x_i, x_l) - d_g(y_j, y_m)|, \quad (3.5)$$

where $d_g(\cdot, \cdot)$ is the normalised geodesic distance between two points on a given surface and $\phi' = \phi \setminus \{(x_i, y_j)\}$ in the most general setting. The optimal bijection ϕ^* being sought minimises \mathcal{D}_{iso} in the space of all $N!$ possible bijections while matching N points. Since a bijection is merely an assignment of a permutation π of the target samples to the fixed source samples, the proposed genetic algorithm efficiently seeks the optimal permutation π^* of indices that will be used as subscripts of $\{y_j\}$, *e.g.*, fixed x_1, x_2, \dots, x_N is assigned to y_4, y_3, \dots, y_{29} , respectively, and $\pi^* = 4, 3, \dots, 29$ (Fig. 3.3).

Having represented a permutation that defines a correspondence as a chromosome, [172] evolves with a fitness function that yields the set of correspondences with minimal distortion using carefully designed genetic operations. The algorithm with the same parameters used in the original paper is able to compute correspondences under articulated, isometric, and non-isometric deformations.

3.5.7 RSDS Wrap 3.3 & 3.4

Wrap 3 [166] is a widely used commercial software tool for mesh processing that includes a wrapping tool that non-rigidly fits one 3D shape to another. It is often

used to wrap a well structured base mesh around unstructured scanned data. This method accepts pre-defined (hard) correspondences between shapes to create control nodes on the source shape. A position for each control node is then found so that it matches the target shape as closely as possible. The method runs iteratively where the density of control nodes is increased per iteration. This leads to an approximation of the target shape with increasing accuracy per iteration.

3.5.8 Partial Functional Correspondence

Rodolà et al. [159] extend the functional mapping framework to address the problem of partial input data. The authors discovered that it is possible to estimate a partial correspondence with functional maps. Given a dense set of point-wise descriptors, computation of a partial functional map \mathbf{C} is formulated as an alternating minimisation problem with two steps. In the first step, correspondences are regularised based on assumptions about the overlap of the source mesh. The slope, orthogonality, and rank of $\mathbf{C}^T \mathbf{C}$ are incorporated to help regularise this step. In the second step, the mapping is regularised by the part. Solutions with a dissimilar area and long boundaries are penalised. This method assumes the deformation to be near-isometric.

3.5.9 Deblurring and Denoising of Maps between Shapes

Ezuz and Ben-Chen [67] observe that functional maps typically lack specificity, moreover recovering vertex-to-vertex correspondences from a functional map leads to undesirable noise. The authors propose a method to refine a functional map, which is capable of recovering vertex-to-point correspondences. This is desirable when transporting high frequency information between surfaces—such as textures. The authors incorporate a correspondence regulariser that favours smooth maps. The method removes noise by first blurring the map, and then applying their proposed smoothness term. Point-wise correspondences are subsequently recovered from the mapping, projected to a facet.

3.5.10 Continuous and Orientation-preserving Correspondences via Functional Maps

Ren et al. [155] propose a method that attempts to address the problem of symmetric ambiguity for functional mapping methods. The method also attempts to ensure that the mapping remains bijective and continuous when recovering point-wise correspondences from a given functional map. **WKS** is used as a point-wise descriptor. It was

shown to be robust to some topological changes in the seminal paper [13], although the descriptor is known to be strictly invariant only under isometric deformation. Unlike most functional map methods, [155] is able to distinguish ambiguous intrinsic symmetries. However, the authors remark that the method is not designed to directly handle partial cases.

3.5.11 FARM+

A variant of [126], referred to here as *FARM+*, is evaluated. This method relies on the registration of a morphable model exploiting the functional map framework [144]. In the original paper, automatic landmarks are detected on protrusions relying only on geometrical information of the *discrete-time evolution process* (DEP) [131] descriptor. These landmarks are used to initialise a dense correspondence in the functional domain, and also to retrieve finer local correspondence over hands and heads. Then, a learned deformable template of human bodies is optimised to fit the target model. Finally, a local refinement is applied to align the template to the target using an **ARAP** regularisation. The correspondence between the template and target is achieved using nearest-neighbour between vertices. This method has been refined by incorporating [125]—removing some iterative steps—and using ZoomOut [133] refinement for functional maps.

3.5.12 CMH Connectivity Transfer

The *coordinate manifold harmonics* (CMH) framework proposed in Marin et al. [127] and extended on animals in Melzi et al. [134] establishes correspondences by transferring the connectivity. The method relies on extending the standard *Laplace-Beltrami operator* (LBO) basis by adding three additional bases that encode extrinsic information of the meshes. This combination of intrinsic and extrinsic information permits the geometry of the models to be fully encoded without information loss due to a low-pass representation. A functional map is then computed, as proposed by Nogneng and Ovsjanikov [142] using a number of landmarks as probe functions. In the case of human bodies, these landmarks are found automatically using DEP [131]. Finally, the connectivity is transferred using the point-to-point correspondence and refined using an **ARAP** regularisation.

3.5.13 ZoomOut

Similarly to CMH [127], *ZoomOut* [133] computes standard LBO bases, which it then refines iteratively. Correspondence and functional map computations are performed alternately, increasing the dimension of the mapping at each step. As with CMH, the connectivity is transferred and the result is refined using **ARAP** optimisation.

3.6 Correspondence datasets

In Tables 3.2 & 3.3 a comprehensive list of datasets are presented. To help differentiate the datasets, a selection of relevant properties have been identified. Rostami et al. [163] provide a survey of existing datasets for evaluating the performance of shape descriptors. While some of the same datasets may be used for both purposes, the focus of this section is datasets that can be used for evaluating correspondences.

Only datasets that have ground-truth correspondences have been mentioned, although, it is possible to assess the performance of correspondence methods on datasets that do not have ground-truths. This can be measured quantitatively by using measures of fitting error (*e.g.*, Hausdorff distance), however this measure is not robust to overfitting. It is also possible to qualitatively present errors using applications that are sensitive to bad correspondences, *e.g.*, texture transfer. The results of qualitative approaches cannot be summarised, so for large collections they are ineffective.

3.7 Non-isometry in datasets

For the problems of non-rigid registration and shape correspondence, there is a great focus on anthropometric—the measurement of humans—surface deformation [157, 176, 8, 202, 78, 21], with the many datasets produced opening up avenues to conduct further research. The popularity of this area is likely due to the numerous pertinent applications that require the accurate measurement of the human body. However, the range of motion of a human is naturally constrained, thus the degree of non-rigid deformation is relatively limited. Whilst significant non-isometric deformation does occur between persons, features like the head, hands and feet, are particularly distinct, enabling correspondences between distinctive surface regions to be established through smooth interpolations. The field of anthropometry may be considered to be a sub-field of morphometry—the measurement of typically organic forms. While morphometry is studied by many other areas of geometry processing (*e.g.*, shape

retrieval, classification, segmentation, etc.), the area is often overlooked by shape correspondence methods as a wider range of deformations must be considered. Notable correspondence datasets for both areas are described here.

The renowned FAUST dataset [21] contains a total of 300 real scans of 10 humans in 30 poses captured using the *3dMDbody.u System* by 3dMD. Subjects were covered in sparse markers to enable shapes to be registered using a novel texture-based technique, which ensured quality alignment in areas with little geometric detail. Ground-truth correspondences between different individuals have been established, with a subset publicly released for training purposes. This was subsequently extended in a dataset that captured human body motion [22], containing 40,000 scans.

Vlasic et al. [202] propose a technique that uses multiple monocular cameras to capture a sequence of images of a human’s performance from multiple angles. At each time step the image from each camera is segmented to separate the background from the actor. This is used to produce a silhouette. A template mesh rigged with a skeleton is used in combination with the silhouettes to reconstruct the human’s pose. This approach enables all models using the same template to have the same connectivity and thus have dense correspondence. 10 performances have been published, in which sequences consist of between 150 and 250 watertight meshes.

The CAESAR dataset [157] is one of the largest human body datasets. 4,431 subjects were scanned in North America, the Netherlands, and Italy using laser scanning. 72 stickers were placed on each subject for use as landmarks; however, due to initial capturing limitations 110 subjects from the European subset do not have landmarks. Subjects were scanned in three poses: standing, sitting comfortably, and sitting with arms raised.

The CAESAR dataset is not publicly available without purchase of a license. The licensing and copyright of content is an issue present across computer vision, especially in research where it may be unclear as to whether one’s work is considered to be for commercial or non-commercial purposes. Whilst still a grey area, subsequent human body datasets have been derived from the dataset [78, 211, 121, 149]. Loper et al. [121] & Zuffi et al. [224] both use data that is not considered to be in the public domain to develop linear blend skinning algorithms to construct a model of the data. These have been subsequently used to produce synthetic datasets for evaluative and training purposes as ground-truths are easily established [74].

Several synthetic datasets have been derived from the Digital Art Zone (Daz Productions, Inc. or Daz 3D). Daz 3D is a digital model platform, as well as a software development company. They have produced a series of *base* (i.e., template)

models which are rigged and include morphs to alter the appearance of a model (*e.g.*, emaciated, muscular, etc.). Many datasets have been derived from these models [32, 158, 149]. Models that use the same template share the same connectivity, and therefore possess dense ground-truths.

Kim et al. [95] combined three existing datasets [8, 32, 111], of which a subset of the shapes are animals, for which a volunteer selected 21 corresponding points on the quadruped shapes. The dataset contains 51 quadruped shapes that approximately contain between 3,000 and 56,000 faces.

Several SHREC tracks have used datasets that contain non-isometric deformation [132, 56, 57, 58]. Melzi et al. [132] and Dyke et al. [58] both curate datasets from existing data sources. These datasets differ in both contents and the processes used to acquire ground-truths. Melzi et al. [132] use an automatic registration method to produce a dense correspondence between 44 human shapes. For [58], a collection of 14 quadruped shapes are manually labelled by three experts. In the case of Dyke et al. [56, 57], scans of inanimate subjects in multiple poses were captured using a hand-held Artec 3D device. Dyke et al. [56] use a small articulated wooden mannequin and wooden hand as base subjects to produce 50 watertight scans using varying textiles to exhibit a range of deformations. Dyke et al. [57] use a stuffed rabbit model made of a stretchy textile with no internal skeleton, enabling greater deformation. The internal filling of the rabbit model is varied with fine and coarse materials that modify the appearance and deformations exhibited.

3.7.1 Dataset taxonomy

Here, the classification criteria for the taxonomy in Tables 3.2 & 3.3 are described.

Deformation type Non-rigid deformations between two shapes may generally be considered to be either *(near-)isometric* or *non-isometric*. During (near-)isometric deformation geodesics are approximately preserved. Conversely, non-isometric deformations involve stretching and compressing that affect geodesics. See discussions in Section 2.4.

Due to scanner limitations when capturing real-world objects, often topological changes may occur (*e.g.*, Fig. 2.13). Such changes increase the complexity of deformations and cause geometric measures, such as geodesics, to become unstable.

For [30], deformations are primarily near-isometric, however there are also some affine transformations.

One notable online source [2] provides a large number of synthetic human bodies that all share the same underlying skeleton rig. A wide variety of animations of these skeletons are available, providing a great potential source for studying action-specific deformation.

Similarity level Many correspondence methods [144, 145, 101, 201, 155] seek to enforce bijectivity between surfaces. Therefore, such methods only perform well on matching problems where the shapes share a *full* correspondence. Other methods are capable of handling injective, or shapes that exhibit a *partial* correspondence. This is a desirable feature, as many real-world correspondence problems involve partial scans.

Scan type Data created either by a computer or an artist using software are considered to be *synthetic*. *Real scans* are representative of real-world scenarios, with realistic deformations, noise and capture limitations. Synthetic datasets can be produced without the use of expensive scanning equipment and may provide incrementally challenging problems, which could be hard to model in real-life. Some synthetic datasets attempt to synthetically induce challenges that are present in particular scanning technologies, *e.g.*, partial shapes [159, 160] and topological changes [102]. A common shortcoming of real scans is that, without either interpolating or automatically registering shapes, establishing a dense ground-truth correspondence between shapes is labour-intensive. In instances where a template shape—either real or synthetic—is deformed, the correspondence is intrinsic to the shape’s triangulation. In the taxonomy, this is referred to as a *deformed template*.

In the case of [21], the evaluative benchmark comprises of real scans, while the training set for learning-based methods provides a deformed template that is watertight.

Ground-truth density A correspondence between two shapes is considered to be *dense* if most vertices in one shape have a corresponding location on the target shape. Otherwise, when the correspondence of only a (relatively small) subset of points is known, the density is *sparse*. In the case of [21, 22] texture markers are registered, while Dyke et al. [56] interpolate sparse texture markers, to produce dense ground-truths.

Ground-truth acquisition As mentioned previously, intrinsic correspondences exist when a collection of shapes all share the same connectivity and therefore have a correspondence. When scanning real-world objects, it is possible to apply structural or texture markers to a surface which are possible to match either *manually* or *automatically* [157, 21, 22, 56, 57]. Where markers are not present, a volunteer can manually establish ground-truth correspondences [95, 58] or—if appropriate—a bespoke shape model [132] can be used.

Training facility For learning-based methods, additional data may be required to train a model for a particular scenario. As the taxonomy demonstrates, such data is scarce. In the case of [56], no training facility was available, but a learning-based method by Groueix et al. [74] used another anthropomorphic linear blend skinning model [121] to produce a training dataset.

Anthropometric Many datasets focus on the problem of deforming human models (*i.e.*, *anthropometry*). Few notable correspondence datasets strictly investigate other problems.

No. of scans The *number of scans* in a dataset is important for the purposes of thorough evaluation. Due to the cost of producing such large datasets, experimental works often contain much fewer scans.

No. of vertices The average *number of vertices* in each mesh is an indicator of the fidelity of the data. The computational cost of present correspondence methods varies widely, with methods that require the computation of dense geodesic distance matrices requiring $\mathcal{O}(n^2)$ space.

A common trend has been to reuse existing datasets, this saves time scanning or creating new digital models, which can take a considerable amount of time. [95, 132, 58] collect models from multiple sources.

As mentioned earlier, few correspondence datasets characterise non-humanoid deformation. This is particularly problematic for learning-based methods—which have been shown to perform particularly well on class-specific deformations—as there are not enough resources to suitably train these methods for additional problems. This is an area that requires urgent focus. Without datasets to facilitate the development of generalised learning-based methods, this area is likely to remain under-researched.

With respect to the work undertaken in this thesis, there is a distinct lack of structure to the deformations in these datasets. Previously, Bogo et al. [21] provided the most structure with the dataset being split into intra-person scan pairs—scans of the same subject with little non-isometry—and inter-person scan pairs—scans of different subjects undergoing significant non-isometric deformation. This problem is addressed from different perspectives by the datasets developed in this thesis in Chapters 5, 6 & 7.

3.8 Evaluative measures

A variety of techniques used to evaluate the performance of correspondence methods on benchmark datasets are discussed by van Kaick et al. [196].

The most popular procedure that is used to assess the performance of correspondence methods is by Kim et al. [95]. The area under curve from this measure is also used as a numerical method for performance comparisons [56, 57, 58].

We shall begin with a few definitions. Let C be a set of correspondences between the source surface X and the target surface Y . For a predicted correspondence between source and target surfaces $c_i = (x_i, y_i) \in X \times Y$, the respective ground-truth correspondence is $(x_i, y_i^*) \in X \times Y$. The distance between the predicted point y_i and the ground-truth point y_i^* is measured using distance function $d_Y(y_i, y_i^*)$, which may represent either the Euclidean or geodesic distance.

Often geodesic distance, or a normalised variant of it, is used to directly measure performance. The average geodesic error may be formulated as

$$E = \frac{1}{|C|} \sum_i^{|C|} d_Y(y_i, y_i^*). \quad (3.6)$$

Many previous evaluative works on shape correspondence have principally used geodesic distance as a measure for error [132, 30, 31, 29]. Note that when measuring the overall error for a collection of target shapes, it may be necessary to normalise the computed error by the surface’s properties such as the area of Y or farthest geodesic on Y .

Measures that require each predicted correspondence to be assigned a binary classification $b(y_i) \in \{0, 1\}$ as either *true positive* (TP) or *true negative* (TN), *e.g.*, precision, recall, and specificity, rely on an appropriate classification strategy $b(y_i) = B(y_i, y_i^*)$. A common strategy is to measure the distance between the predicted and ground-truth correspondence points $d_Y(y_i, y_i^*)$; points that are below a specified error

threshold ϵ are considered correct. This description is the basis of the correspondence error measure proposed by Kim et al. [95], where a correspondence is considered to be a TP when $d_Y(y_i, y_i^*) \leq \epsilon$. The value of ϵ is increased to measure the number of TP over larger radii, which can be used to produce a curve. This approach fails to characterise the distribution and sparsity of correspondences on shapes where a limited number of ground-truth correspondences are available.

The benchmark protocol described by Kim et al. [95] is considered to be the standard error measure for correspondences. The normalised geodesic error may be used to produce further statistics through the use of **AUC**, used in [56]. For a collection of shapes, Kim et al. [95] also report the average of the maximal geodesic error. Rodolà et al. [159] & Cosmo et al. [43] report the average geodesic error over a dataset of shapes with gradually reducing surface areas to measure the robustness of methods on increasingly partial scans. For functional mapping approaches, Corman et al. [42] measure the quality of ground-truth and predicted functional basis. In the case of registration methods, where one shape is deformed to align with another, it is possible to measure fitting error using the Hausdorff distance [40, 10]. A variety of other validation methods for shape correspondences are discussed by van Kaick et al. [196].

Qualitative techniques using visual mappings between two shapes in which topological information is transported (*e.g.*, colour or texture transfer) [66, 60] are also used. However, these techniques are not an effective way to succinctly summarise the performance of a method on larger datasets or for comparing the performance of multiple methods. Further evaluation may be performed on a proposed registration or correspondence algorithm using an application that requires a correspondence mapping (*e.g.*, shape retrieval [33], consistent quadrangulation [66, 14, 143]).

For evaluating shape descriptors there are several other useful evaluation measures, Guo et al. [76] discuss a few. *Receiver operating characteristic curve* (**ROC**) and *cumulative match characteristic* (**CMC**) are often used to help evaluate and compare performance—Moon and Phillips [137] discuss these measures further. These descriptors may be used to initialise typical shape correspondence pipelines, and can be subsequently evaluated using shape correspondence tools. To demonstrate the quality of their feature descriptor without influence from the appended correspondence method, Masci et al. [128] simply use nearest neighbours in the descriptor space.

Despite the wealth of evaluative measures available, there is still scope for further development in this area. Through the evaluative experiments undertaken in Chapter 5, it is discovered that the benchmark protocol by Kim et al. [95] does not report

information such as sparsity. Chapter 7 addresses this through the use of a novel evaluative measure that seeks to help characterise both the sparsity and distribution of correspondences.

3.9 Medical image registration of skulls

Chapter 8 investigates the problem of registering the orbits of skulls scanned using *computerised tomography* (CT). By registering the skulls, geometric measures can be calculated automatically. This is critical in scenarios where manual labelling is infeasible, such as for large datasets. While specifically automatic orbit registration is an unstudied area, the general registration of medical images is a well-studied problem. In this section, only key existing skull registration pipelines are discussed.

Utkualp and Ercan [195] provide a broad overview of the development of anthropometric techniques in the field of medicine. With the invention of *magnetic resonance imaging* (MRI) and CT devices, the field of medical image registration was quickly established. In medical image registration, numerous approaches have been proposed for a variety of tasks. Many methods are designed to register images of the same subject—this only requires a rigid transformation to be estimated. However, in practice, global affine-based registration approaches are typically used to accommodate scanner calibration errors [46]. For works that operate directly on the original volume of inter-subject scans, a non-rigid B-spline registration technique is often used [180].

For the work in Chapter 8, a triangulated mesh representation of the skull’s surface is extracted for analysis. Few works have considered the problem of inter-subject skull registration in this form [51, 218, 174, 219].

Deng et al. [51] seek to estimate exterior facial geometry. First, a template skull is coarsely registering to a target skull using *thin-plate splines* (TPS). Then, the deformation is locally refined using radial basis functions that support local non-isometry, minimising global distortion. The deformation is then applied to the known exterior facial geometry of the template to reconstruct the face. This method relies on manually placed landmarks, therefore making it infeasible to apply to larger databases.

Zhang et al. [218] automatically establish landmarks between extracted skulls in a coarse-to-fine manner. First, global affine-based registration is used to estimate the location of a set of sparse landmarks. Subsequently, the locations are iteratively refined in an alternating manner to optimise the location of points on the mid-sagittal plane, and the Frankfort horizontal plane. The resulting registration is then refined

non-rigidly using a **TPS**-based approach. This method relies on consistent landmark detection, which is difficult to guarantee when processing low-quality data.

Sahillioglu and Kavan [174] propose a method to register skulls with significant deformities. A multi-modal approach is used, with a triangulated source mesh being fitted to a target volume comprising of voxels. The method incorporates *principal component analysis* (PCA), and therefore has a low level of computational complexity. In Chapter 8, the proposed approach differs, as a triangulated mesh representation is used for both the source and target.

Zhao et al. [219] propose a conformal parameterisation method for registering skulls captured using **CT**. This method is shown to work on simple skulls with holes, but only one wall. To extract only the exterior skull bone information, the image is discretised using marching cubes, and then internal geometry is removed. Although it might be possible to reformulate the extraction method to cut more conservatively, this technique can remove necessary information that is on, or near, the orbit.

Pinheiro et al. [150] present a pipeline that is designed to measure craniofacial symmetries between a series of patient scans captured over a number of years. While related, this work only considers scans of the same individual. The authors demonstrate the performance of the method on a set of synthetic deformations of a single shape—therefore bypassing challenging geometric errors that occur in real scans. Furthermore, this work measures asymmetry globally over the entire skull, whereas, the work in Chapter 8 seeks to only measure asymmetry at the pair of orbits.

Other works have investigated the use of adaptive templates to improve the subsequent surface registration (*e.g.*, [48]). Dai et al. [47] construct a statistical shape model, which decomposes the symmetric and asymmetric components of a dataset of heads. A large collection of scans of human faces with textures were used. Landmarks were automatically found in each texture using a supervised 2D localisation approach [222]. If an automatic localisation method is not available, such an approach is not be feasible without extensive additional manual labelling.

Other registration-based applications have been explored. There has been a great focus on combining data from multiple sources [79, 183, 23]. These have been shown to help automate parts of tasks, such as surgery planning [79], and forensic reconstruction [198], which would typically be performed manually.

Overall, there are several undesirable problems with state-of-the-art methods for skull registration. For large datasets with significant variations, it can be infeasible to obtain high-quality landmarks (*e.g.*, [51, 218]). **CT** scans can contain problematic noise, which is exacerbated when processing low-quality data. Therefore, methods

that are sensitive to topological and geometric error are not suitable for this problem (*e.g.*, [219]). Additionally, although not addressed in this thesis, an additional problem is the absence of studies that evaluate these methods in the context of intra-subject skull registration.

3.10 Contributions

This chapter has explored current notable works in the area of non-rigid registration and its related fields. The area of registration is still maturing, with many avenues of research yet to be explored or requiring further attention.

For shape correspondence methods, there has been a significant amount of work that considers isometric deformation. Recently, a growing focus has been on non-isometric deformation [201, 200, 67, 155, 66]. Unfortunately, the assumptions made in previous works are incompatible with non-isometric deformation; therefore, new techniques are required for handling challenging cases such as partial correspondence.

Many works focus their evaluations on synthetic benchmark datasets. With respect to real-world scenarios, these datasets do not appropriately reflect the challenges presented in real scans. This is a persistent problem in many works, where the quality of the data used differs from the envisioned application. The inclusion of thorough a evaluation is necessary for potential users to discern whether a method is applicable for a particular scenario or capture modality.

Existing datasets indicate a dominant interest in anthropometry. It should be noted that there are other reasons why it is much easier to capture human data than other objects or species (*e.g.*, copyright, size, compliance, etc.). These are considered in Chapter 7 where a dataset is curated that consists of 3D scans of quadrupeds that are available on the internet.

3.11 Summary

The aim of this chapter was to provide a comprehensive overview of the key literature within the scope outlined in Section 3.1. The amount of literature on non-rigid registration is substantial, and it is not possible to cover all of it here. However, this chapter has examined the literature that is most relevant to the problems investigated in this thesis. Hereon, little additional literature is introduced. For the benchmark chapters (Chapters 5, 6 & 7), the reader is directed to the descriptions of the methods in Section 3.5.

Table 3.1: A taxonomy of the methods evaluated in Chapters 5, 6 & 7, complementing the survey of Sahillioglu [171].

Method	Criteria										
	Similarity level		Deformation type			Shape processing			Output density		
	Full	Partial	Isometric	Non-isometric	Pairwise	Collectionwise	Sparse	Dense			
R3DS Wrap 3	✓		✓		✓						✓
Ezuz and Ben-Chen [67]	✓				✓						✓
Rodolà et al. [159]	✓		✓								✓
Ren et al. [155]	✓				✓						✓
Bouaziz and Pauly [27]	✓		✓		✓						✓
Li et al. [109]	✓										✓
Vestner et al. [200]					✓						✓
Melzi et al. [133]	✓				✓						✓
Marin et al. [127]	✓				✓						✓
Sahillioglu [172]	✓		✓								✓
Groueix et al. [74]	✓				✓						✓

Method	Criteria continued										
	Solution approach				Speed			Surface topology			
	Registration	Similarity	Yes	No	Automatic	Semi	Fast	Med	Slow	Arbitrary	Sphere
R3DS Wrap 3	✓		✓	✓	✓		✓				✓
Ezuz and Ben-Chen [67]			✓	✓	✓		✓				✓
Rodolà et al. [159]			✓	✓	✓		✓				✓
Ren et al. [155]			✓	✓	✓		✓				✓
Bouaziz and Pauly [27]	✓		✓	✓	✓		✓				✓
Li et al. [109]	✓		✓	✓	✓		✓				✓
Vestner et al. [200]			✓	✓	✓		✓				✓
Dyke et al. [55]			✓	✓	✓		✓				✓
Melzi et al. [133]			✓	✓	✓		✓				✓
Marin et al. [127]			✓	✓	✓		✓				✓
Sahillioglu [172]			✓	✓	✓		✓				✓
Groueix et al. [74]			✓	✓	✓		✓				✓

Table 3.2: A taxonomy of benchmark datasets for evaluating correspondence methods.

Method	Deformation type		Topological change	Similarity level		Scan type			Ground-truths	
	Near-isometric	Non-isometric		Partial	Real	Synthetic	Deformed template	Density Sparse	Dense	
[157]		✓	✓	✓					✓	
[187]	✓				✓		✓			
[8]	✓						✓			
[32] hi-res	✓						✓			✓
[32] lo-res	✓						✓			✓
[202]	✓				✓		✓			✓
[78]		✓			✓		✓			✓
[30]	✓*			✓			✓			✓
[95]		✓			✓		✓			✓
[21]		✓		✓			✓*			✓*
[158]	✓						✓			✓
[211]							✓			✓
[102]	✓		✓				✓			✓
[159, 43]	✓						✓			✓
[160]							✓			✓*
[22]		✓	✓	✓			✓			✓
[151]		✓					✓			✓
[56]		✓	✓	✓			✓			✓
[132]		✓	✓	✓			✓			✓
[2]	✓									✓
[57]		✓	✓		✓		✓			✓
[58]		✓	✓		✓		✓			✓

Table 3.3: Correspondence dataset taxonomy continued.

Method	Ground-truths (cont.)			Training facility	Anthropometric	Other	No. of scans	No. of vertices	
	Manual	Automatic	Intrinsic					$\leq 25,000$	$> 25,000$
[157]	✓		✓		✓	4431*	✓	✓	
[187]			✓		partially	551		✓	
[8]			✓		✓	71		✓	
[32] hi-res			✓		partially	80		✓	
[32] lo-res			✓		partially	148	✓		
[202]			✓		✓	1500	✓		
[78]	✓	✓			✓	550		✓	
[30]			✓		partially	56		✓	
[95]	✓		✓*		partially	551	✓	✓	
[21]		✓		✓	✓	300		✓	
[158]			✓		partially	32		✓	
[211]		✓			✓	3000	✓	✓	
[102]			✓		✓	26		✓	
[159]			✓		partially	1216		✓	
[160]			✓	✓	partially	4258		✓	
[22]		✓		✓	partially	40000		✓	
[151]		✓			✓	4309	✓	✓	
[56]	✓				✓	50		✓	
[132]		✓			✓	44		✓	
[2]	✓				✓	121 (base models)		✓	
[57]	✓					12		✓	
[58]	✓					14		✓	

Table 3.4: A list of datasets that have been reused by other benchmark datasets.

Original dataset	Derivative datasets
Robinette et al. [157]	[78, 211, 121, 151, 132]
Anguelov et al. [8]	[132]
Bronstein et al. [32]	[95, 159, 43, 160, 132]
Vlasic et al. [202]	[95]
Lian et al. [111]	[95]
Yang et al. [211]	[132]
Bogo et al. [21]	[132]
Rodolà et al. [158]	[102]

Chapter 4

Non-rigid registration under anisotropic deformations

Overview

This chapter introduces a novel non-rigid registration method, which is capable of handling anisotropic deformation. Section 4.1 describes the problem and identifies the key contributions of this method. Relevant literature is discussed in the context of this chapter in Section 4.2, this extends the literature reviewed in Chapter 3. Section 4.3 details the novel components of the proposed pipeline. The performance of the proposed method is compared with other state-of-the-art methods on several notable real and synthetic benchmark datasets through sets of quantitative and qualitative results. Further evaluation of this method is undertaken in Chapters 5, 6 & 7. Key limitations for the proposed method are mentioned in Section 4.5. Finally, some concluding thoughts and suggestions for future work are given in Section 4.6.

4.1 Introduction

Surface registration is a fundamental problem in the domains of computer graphics and vision, in which the aim is to find a transformation that best aligns two input surfaces. Surface registration algorithms underline computational solutions to many prevailing problems, such as 3D acquisition/reconstruction, statistical shape analysis and shape retrieval. With the increasing pervasiveness of 3D scanners and application of 3D scanning in real-world scenarios, the importance of accurate registration algorithms is continuing to rise.

In many real-life scenarios, surfaces are often non-rigidly deformed. Non-rigid surface registration is therefore required to find the non-rigid transformation between

them. Extending from the well-known *Iterative Closest Point* (**ICP**) approach for rigid registration [19, 39], *Non-rigid ICP* (**N-ICP**) methods [27] achieve registration for non-rigid surfaces by alternating between two steps. In the first step, a set of correspondences is computed using a closest point criterion, and then the second step identifies a non-rigid transformation that minimises an error metric. Generally, the associated cost function decreases after each iteration, converging monotonically to a local minimum. Because of the simple way correspondences are generated, **N-ICP** is fast enough to be used in some real-time applications; though alone, it is incapable of coping with large-scale deformations. **N-ICP** methods thus typically require an initial set of correspondences generated through alternative means, such as automatic/manual markers (*e.g.*, [6]), so as to achieve good registration results.

In the **N-ICP** framework, dense correspondences are obtained using the closest point criterion. However, it is only effective when two surfaces are reasonably close. Additional (sparse) correspondences are often needed to cope with shapes with large deformations. Since shapes to be registered may only have partial overlaps (*e.g.*, due to occlusion when two views are captured), correspondences are most often generated by feature matching of local shape descriptors of source and target shapes.

This requirement exhibits two challenges: First, it can lead to false matches between areas that appear the same locally, but belong to different regions (*e.g.*, any local regions on a cylindrical surface). This could be accounted for by ensuring the correspondences are consistent. Second, in cases where the surface has been warped, they may not match (unless the shape signatures are insensitive to the particular warp, which is not generally possible). Typically, there is little contingency built into these **N-ICP** methods, with most relying on such areas being sufficiently insignificant so as to not affect the final result.

Most recent correspondence methods (*e.g.*, [82, 192]) used in non-rigid registration account for a certain degree of near-isometric deformation by employing a geodesic-based consistency measure. However, for surfaces with large deformations, non-isometric deformations commonly exist (typically in joint areas of articulation, or on deforming parts), making the global isometry assumption less useful. In the literature, most surface registration methods do not directly address non-isometric deformation. They simply penalise any deformation that is non-isometric. To cope with large deformations, a new consistency measure that takes into account anisotropic non-isometric deformation explicitly is proposed.

Some previous work addresses a related problem of establishing correspondences between shapes which involve potentially large deformations [95, 117]. However,

fundamental differences exist. Non-rigid registration aims to identify non-rigid deformation (usually in the form of a set of local deformations) to deform the source mesh to align with the target, whereas correspondence methods only identify point-to-point correspondences between surfaces. On the one hand, in order to directly derive deformation with the latter approach, complete or dense per-vertex correspondences would be needed. The computation of such dense correspondences can be slow. When two surfaces are only partially overlapping, such complete correspondences may not be defined. On the other hand, these methods may have significant restrictions on topology (complete genus-zero models) [95] or may generate globally inconsistent correspondences due to the lack of a global constraint [117]. Compared to these methods, non-rigid registration can often be driven by a sparse set of reliable correspondences which is easier and more efficient to obtain. It also does not require any strong topology assumptions.

Due to the typical sparsity of the correspondences used, non-rigid registration usually requires some regularisation to ensure that the surface deforms appropriately during the **N-ICP** iterations. Notable methods define the local regularisation neighbourhood to be one-ring (*i.e.*, vertices connected to the vertex of interest by an edge) [27] using an as-rigid-as-possible (ARAP) formulation [185]. This however is insufficient for cases with large deformations, as shown in our experiments. Following work by Chen et al. [38], which uses r -ring **ARAP** energy for deformation with controllable stiffness, r -ring **ARAP** regularisation is incorporated to cover larger neighbourhoods, allowing large deformation to be handled effectively in the **N-ICP** framework.

Technical Contribution In this chapter, a novel non-rigid registration technique capable of registering large-scale and non-isometric deformations is proposed. Addressing large-scale and non-isometric deformation is a fundamental challenge of existing non-rigid registration techniques. To the best of our knowledge, this challenge has not been addressed before. More specifically, our main technical contributions are:

- A novel method to estimate anisotropic deformations on a discrete mesh by using the principal scaling factor.
- A correspondence generation and correspondence pruning method based on local geodesics that copes with anisotropic deformations; this makes use of our anisotropic deformation estimate.

- The r -ring **ARAP** formulation is introduced for regularisation in non-rigid registration, which effectively handles challenging large deformations where existing registration methods fail.

Both qualitative and quantitative evaluation are performed using public benchmark datasets. Our results show that our method outperforms the state-of-the-art methods in non-rigid registration.

4.2 Related work

Here we briefly overview methods which are particularly related to this chapter. Please refer to Chapter 3 for a more extensive literature review related to this thesis.

Shape correspondences A variety of methods have been proposed for establishing correspondences between shapes. Notable works have been discussed in Chapter 3.

For non-rigid registration, it is usually sufficient to have a sparse set of initial correspondences. Leordeanu and Hebert [104] first used spectral matching to acquire a consistent set of correspondences between images. Huang et al. [82] adapt the idea to address large-scale, but piecewise-rigid deformations. They further use an assumption of global isometry to obtain correspondences. However, the computation of geodesic distances is slow, and using an isometry measure based on global geodesic distances may not adequately model real-life deformation well. Due to these drawbacks, Tam et al. [192] propose to use a more flexible local isometry global consistency model, along with a more efficient diffusion framework to select reliable correspondences. All these methods assume (near-)isometric deformation. However, these techniques still cannot model real-life deformation well. Here, to the best of our knowledge, this is the first work to explicitly model anisotropic non-isometric deformation in a pruning and registration technique, allowing reliable correspondences and non-rigid alignment to be obtained even under large-scale deformation.

Shape deformation A detailed discussion of deformation regularisation techniques is given in Section 3.3.3. The regularisation term restricts arbitrary deformations, and ensures a natural deformation. Previous work such as Achenbach et al. [1], only uses anisotropic regularisation in the later refinement stage and therefore the method is not designed for large-scale anisotropic deformations, which is addressed in this chapter.

As mentioned in the previous chapter, Chen et al. [38] propose a simple formulation that extends [185], enlarging the neighbourhood size (to r -ring neighbours) to change

the local rigidity of deformations. In this work, the r -ring **ARAP** energy is introduced to the **N-ICP** framework and we demonstrate its effectiveness in handling large-scale deformations. This approach may seem counter-intuitive, as it penalises anisotropic deformation. However, when it comes to challenging large-scale deformations where reliable correspondences may not be readily available, such regularisation is necessary to ensure the deformation is not arbitrary and does not overfit to noise.

Non-rigid registration Non-rigid registration aims to align surfaces undergoing non-rigid deformation. Much research has been conducted due to its general applicability, see Sections 3.3 & 3.9. However, most research makes strong assumptions regarding the type of deformation, *e.g.*, piecewise-rigid [82] and isometric [210, 109]. Although these assumptions are reasonable when the deformations are small or near-isometric, it is still a challenge to register surfaces with large-scale non-isometric deformation, in which case anisotropic deformation is common. Jiang et al. [88] recently proposed a method that incorporates a regularisation model that is capable of handling non-isometric deformations. However, the method relies on a farthest point sampling technique to establish coarse correspondences, which can be sensitive to topological change.

4.3 Methodology

Our surface registration algorithm (Algorithm 1) consists of four key components: modelling of anisotropic non-isometric deformation (Section 4.3.2), correspondence computation (Section 4.3.3), correspondence pruning (Section 4.3.4) and deformation optimisation (Section 4.3.5). These components are underpinned by our incorporation of an anisotropic geodesic method. In this section, we first present an overview of the algorithm, and subsequently detail each component in the following subsections.

4.3.1 Algorithm overview

Given a source mesh X and target mesh Y , geodesic distances between vertices are pre-computed on both meshes (d_X, d_Y). Note, at this stage, no assumptions about non-isometry are made, so the distance measured between two points should approximate the conventional geodesic distance (as illustrated in Fig. 4.1).




Figure 4.1: An isotropic geodesic.

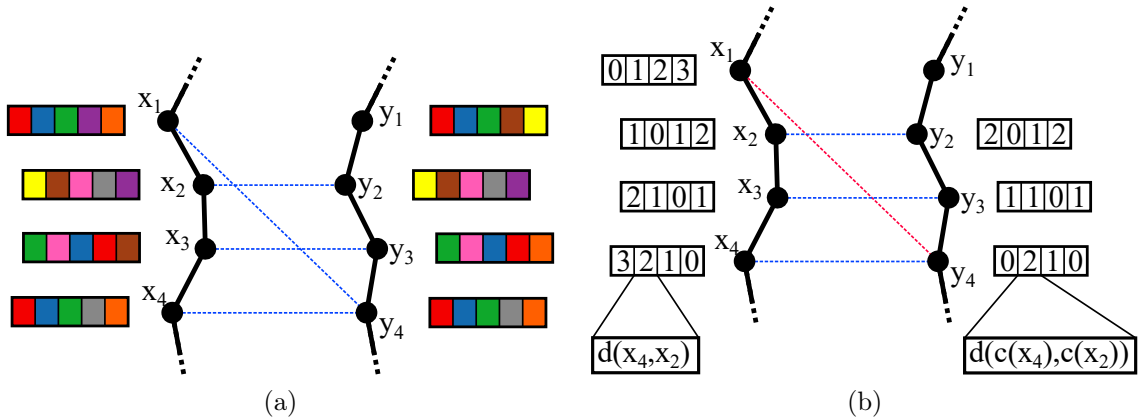


Figure 4.2: (a) Closest matching descriptor between X and Y , the groups of coloured squares represent feature vectors. (b) Pruning descriptors based on geodesic consistency, the boxes with integer values represent the distance to neighbouring points in X and the corresponding points in Y .

Next, shape descriptors are computed at each vertex.

Since identifying the optimal feature descriptor for non-rigid surfaces is not the main focus of this work, **SHOT** [194] is used to establish initial correspondences due to its robust performance. Pairs of vertices between shapes are matched based on their similarity (see Fig 4.2a). The top k matches may be found efficiently using nearest neighbour distance in the descriptor space using a k -d tree.

Each pair of matched points share a similar a feature description; however, if a similar feature is present in multiple regions, these matched pair might not share the corresponding location. The pruning step—on line 4—seeks to address this by removing matches where geodesic distances to neighbouring points are inconsistent between surfaces (see Fig. 4.2b). Still, no assumptions about non-isometry are made, so the original pre-computed geodesics (d_X, d_Y) are used here.

An extended **N-ICP** algorithm is then used to non-rigidly align the two surfaces, initially using the pruned matches as hard correspondences to coarsely align stable points. An illustration of this step is shown in Fig. 4.3. The outputs of line 5 are the deformed mesh X' and a dense set of nearest-neighbour correspondences C_{N-ICP} .

Up to this point, the described pipeline is suitable for (near-)isometrically deforming surfaces. The components of the method described hereon seek to handle non-isometries. On line 7, given a dense correspondence between X & Y , the eigenvectors/values that correspond to the scaling factor in the principal scaling direction and orthogonal scaling direction are computed. In each iteration of the pipeline, the




Figure 4.3: **N-ICP** using a coarse set of hard correspondences. To the left, the source surface (X), represented by the red line, is non-rigidly aligned to the target surface (Y), in blue. This produces the deformed surface X' , in red, on the right.

original geodesics (d_X, d_Y) are incorporated. The pairs of eigenvectors/values form an anisotropic tensor field, shown in Fig. 4.4.

This information is used to re-scale (or weight) the edge lengths of each triangle with respect to the local field. This mesh with modified edge information is used to compute anisotropic geodesics d_X^A using the method proposed by Liu et al. [119]. Fig. 4.5 shows an example of the optimal path between two points after applying the anisotropic tensor field.




Figure 4.5: An anisotropic geodesic.

The dense set of correspondences C_{N-ICP} are pruned to discard pairs of correspondences with inconsistent geodesics to neighbouring points—these typically occur in misaligned regions. In contrast to line 4, rather than using the geodesics that were initially computed (d_X), the anisotropic geodesics (d_X^A) are used for the metric on X .

The pruning stage often discards correspondences between regions that are misaligned. The *correspondence inference* stage re-populates these areas with correspondences by locating the most geodesically compatible point on the target surface using the remaining unpruned correspondences (C_{pruned} from line 9). This concept is illustrated in Fig. 4.6.

The correspondences from the pruning step on line 9 (C_{pruned}) and the correspondence inference step on line 10 ($C_{inferred}$) are combined into a single set. The




Figure 4.4: The resulting pairs of eigenvectors/values of each corresponding point between X and Y based on the anisotropic non-isometric stretch.




Figure 4.6: Given the point x_2 on surface X , the geodesic compatibility of correspondences is measured for each candidate correspondence on Y (i.e., y_2 & y_3). In this case, the distance values between y_2 and the known correspondences are more consistent with the distance values of x_2 , than the values of y_3 .

composite set is then pruned. Since the geodesics between the correspondences in C_{pruned} are already consistent, they are unlikely to be removed during the pruning step. However, C_{inferred} may introduce correspondences that are not consistent with neighbouring candidate correspondences in the C_{inferred} set. The poorer correspondences in C_{inferred} may not help in the alignment process; therefore, they are removed in this pruning step.

Finally, the new C_{pruned} is used to help initialise the extended **N-ICP** algorithm. This final pruning stage is necessary as some correspondences introduced in the correspondence inference stage may not be adequately consistent. If further iterations are required, the $C_{\text{N-ICP}}$ produced in this step is next used to compute the principal scaling factor.

Algorithm 1: Algorithm overview

Input : X, Y – Source and target surfaces

Output: X' – Deformed surface (X')

$C_{\text{N-ICP}}$ – correspondences

```

1  $d_X \leftarrow \text{ExactGeodesics}(X)$ 
2  $d_Y \leftarrow \text{ExactGeodesics}(Y)$ 
3  $C_{\text{FeatureMatch}} \leftarrow \text{FeatureMatch}(X, Y)$  // Local feature matching.
4  $C_{\text{pruned}} \leftarrow \text{prune}(d_X, d_Y, C_{\text{FeatureMatch}})$ 
5  $X', C_{\text{N-ICP}} \leftarrow \text{N-ICP}(X, Y, C_{\text{pruned}})$ 
6 while  $X'$  has sufficient changes over the last iteration do
7    $E_X \leftarrow \text{PrincipalScalingFactor}(X, Y, d_X, d_Y, C_{\text{N-ICP}})$  // Output source edge
   lengths.
8    $d_X^A \leftarrow \text{AnisotropicGeodesics}(X, E_X)$ 
9    $C_{\text{pruned}} \leftarrow \text{prune}(d_X^A, d_Y, C_{\text{N-ICP}})$ 
10   $C_{\text{inferred}} \leftarrow \text{CorrespondenceInference}(X, Y, d_X^A, d_Y, C_{\text{pruned}})$ 
11   $C_{\text{pruned}} \leftarrow \text{prune}(d_X^A, d_Y, C_{\text{inferred}} \cup C_{\text{pruned}})$  // Use scaling factor in
   consistency measure.
12   $X', C_{\text{N-ICP}} \leftarrow \text{N-ICP}(X, Y, C_{\text{pruned}})$ 
13 end

```

4.3.2 Anisotropy estimation

In real-life non-rigid scenarios, anisotropic non-isometric deformation is common. For example, the contraction of the muscle bicep would cause the arm to bend at the elbow (as shown in Fig. 4.16). Such muscle contraction involves shortening in longitudinal and lengthening in lateral directions of the muscle. It is a typical anisotropic non-isometric deformation. While some existing works incorporate regularisation terms

that do allow for some non-isometry [6, 138], alone, the data terms used are only sufficient to handle modest amounts of deformation. Therefore, additional correspondence information is necessary to align challenging regions of a shape.

To better model these anisotropic non-isometric deformations for a registration technique, we develop a local anisotropy metric, and incorporate it in the local geodesic computation. This allows more and better correspondences to be found in non-isometrically deforming areas, supporting more robust registration results. Specifically, the anisotropic measure is incorporated into the estimate of geodesic distance, which is used to accurately measure the consistency of pairs of correspondences. In Algorithm 1, the anisotropic measure is used in steps 9, 10 & 11. In the following, we first discuss the anisotropy metric in Section 4.3.2.1, then discuss how to adopt the metric for geodesic computation in Section 4.3.2.2.

4.3.2.1 Local anisotropic deformation metric To model the local anisotropic deformation at each vertex, we propose an anisotropy metric—a pair of local eigenvectors that defines the *Principal Scaling Factor* (PSF) of deformation.

We first project local points in the source mesh onto the tangent plane of the vertex of concern. Then we can estimate the local scaling factor by comparing these 2D coordinates with the corresponding geodesic distances from the target mesh (illustrated in Fig. 4.7a). This allows us to extract the scaling directions and magnitudes via eigen-analysis, inspired by the idea of principal curvatures [64]. These constitute one part of a 2-order tensor field that describes changes in scaling at a per-vertex level of the source shape.

In order to estimate PSF, it is necessary to have some points with known correspondences. This leads to a chicken-and-egg problem. We address this issue with an iterative approach. We first compute initial correspondences by matching features between shapes, as per step 3 of Algorithm 1. Consider a discrete vertex \mathbf{s}_i on the source shape X , with neighbouring vertices $\mathbf{s}_j \in N_X(\mathbf{s}_i)$ where the correspondences of \mathbf{s}_i and all \mathbf{s}_j are known. We project \mathbf{s}_i onto a 2D plane with an arbitrary orientation that is tangential to the normal of \mathbf{s}_i . Let $\hat{\mathbf{s}}_i$ and $\hat{\mathbf{s}}_j$ be the projection of vertices \mathbf{s}_i and \mathbf{s}_j onto the 2D plane, respectively. The corresponding points on the target shape for \mathbf{s}_i and \mathbf{s}_j on the source shape are denoted by \mathbf{t}_u and \mathbf{t}_v .

For simpler formulation, we set $\hat{\mathbf{s}}_i = (0, 0)$ as the origin, and denote $\hat{\mathbf{s}}_j = (\hat{x}_j, \hat{y}_j)$. We also define $d_j = d_g(\mathbf{t}_u, \mathbf{t}_v)$ as the geodesic distance on the target shape between \mathbf{t}_u and \mathbf{t}_v . To ensure robustness, we formulate the problem as a least squares fitting, and use a quadratic function $f(\hat{x}, \hat{y})$ to approximate d_j^2 since low-order polynomials

are more stable. The squared distance should have its centre of symmetry at the origin and $f(0,0) = 0$, at which the x term, y term and constant thus should all be 0, leading to

$$f(\hat{x}, \hat{y}) = a\hat{x}^2 + b\hat{x}\hat{y} + c\hat{y}^2, \quad (4.1)$$

where a , b and c are coefficients to be determined. We formulate the least squares problem to minimise the following

$$F(a, b, c) = \sum_{\mathbf{s}_j \in N_X(\mathbf{s}_i)} (a\hat{x}_j^2 + b\hat{x}_j\hat{y}_j + c\hat{y}_j^2 - d_j^2)^2. \quad (4.2)$$

Setting $\frac{\partial F}{\partial(a,b,c)} = 0$ leads to a linear system, which can be easily solved. Let $\hat{\mathbf{x}} = (\hat{x}, \hat{y})$. We construct a matrix \mathbf{A}_q

$$\mathbf{A}_q = \begin{bmatrix} a & \frac{1}{2}b \\ \frac{1}{2}b & c \end{bmatrix}, \quad (4.3)$$

and $f(\hat{\mathbf{x}}) = \hat{\mathbf{x}}^T \mathbf{A}_q \hat{\mathbf{x}}$. We then find the first and second eigenvectors of \mathbf{A}_q , *i.e.*, $\hat{\mathbf{e}}_1$, $\hat{\mathbf{e}}_2$ and their respective eigenvalues ($\|\hat{\mathbf{e}}_1\|$, $\|\hat{\mathbf{e}}_2\|$). These eigenvectors represent the anisotropy metric at a specific vertex. We may calculate $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{R}^3$ by mapping the tangent plane that $\mathbf{e}_1, \mathbf{e}_2$ lie on back into \mathbb{R}^3 . This process is repeated for each vertex in X with a correspondence on Y .

4.3.2.2 Measuring anisotropic geodesic distance To measure geodesic distance we use the method proposed by Liu et al. [119]. This method requires the anisotropy metric to be encoded into the predicted edge lengths of the shape. We update the estimated lengths of the edge set for the source shape using the following procedure.

Let \mathbf{v}_{ab} be a vector that represents an edge between two adjacent vertices on the source shape (*i.e.*, $\mathbf{v}_{ab} = \mathbf{s}_b - \mathbf{s}_a$). At each vertex, there is a pair of eigenvectors, a principal and secondary vector ($\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2$). Thus, for each edge, there are two pairs of eigenvectors ($\hat{\mathbf{e}}_{1a}, \hat{\mathbf{e}}_{2a}$) & ($\hat{\mathbf{e}}_{1b}, \hat{\mathbf{e}}_{2b}$). The eigenvalue is encoded into the length of each eigenvector $\|\hat{\mathbf{e}}_\alpha\|$. An example of this in the 2D tangent plane is illustrated in Fig. 4.7b.

In practice, \mathbf{e}_α and \mathbf{v}_{ab} are embedded in \mathbb{R}^3 . We apply the following calculation for both points a and b . For simplicity, we only describe the process for one point. We assign the normal vector at the point of concern to $\mathbf{e}_3 = \frac{\mathbf{e}_1 \times \mathbf{e}_2}{\|\mathbf{e}_1 \times \mathbf{e}_2\|}$. We use the three normalised vectors as the current basis (*i.e.*, $\mathbf{B} = (\frac{\mathbf{e}_1}{\|\mathbf{e}_1\|}, \frac{\mathbf{e}_2}{\|\mathbf{e}_2\|}, \mathbf{e}_3)$, as \mathbf{e}_3 is already normalised). Transforming the subject edge into this basis ($\tilde{\mathbf{v}}_{ab} = \mathbf{B}^{-1}\mathbf{v}_{ab}$), the components of each axis now represent the amount of stretching in each direction

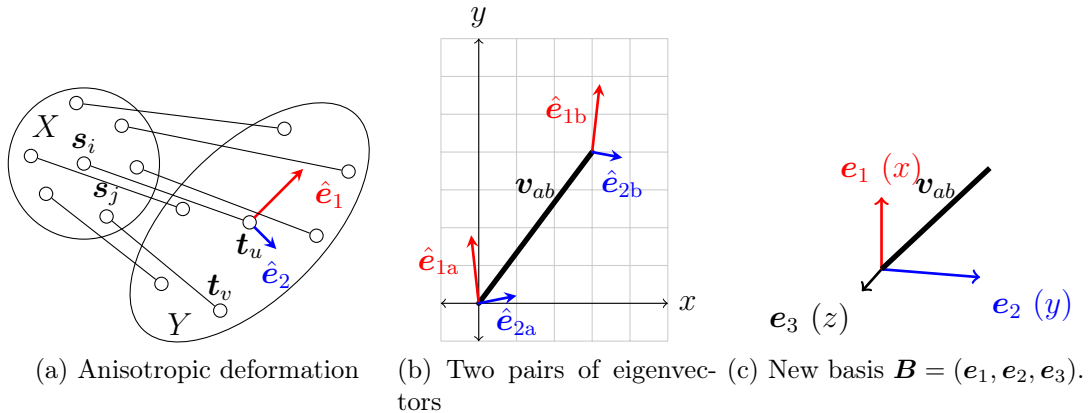


Figure 4.7: Illustrations that demonstrate the idea of estimating anisotropic non-isometric deformation. (a) An anisotropic deformation of points between two shapes (X & Y) may be approximately modelled by an ellipse. The surrounding unlabelled points represent neighbouring points on the source shape $\mathbf{s}_j \in N_X(\mathbf{s}_i)$ and the predicted correspondence of \mathbf{t}_v on surface Y . (b) Two pairs of eigenvectors ($\hat{\mathbf{e}}_{1a}, \hat{\mathbf{e}}_{2a}, \hat{\mathbf{e}}_{1b}, \hat{\mathbf{e}}_{2b} \in \mathbb{R}^2$) represent stretching for an edge \mathbf{v}_{ab} . (c) \mathbf{v}_{ab} transforming into new basis $\mathbf{B} = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$.

(n.b. \mathbf{e}_3 always has a stretching factor of 1). We illustrate this new basis w.r.t. point a in 3D in Fig. 4.7c. We can then assign each of our scaling factors to an axis as follows: $x : \|\mathbf{e}_1\|$, $y : \|\mathbf{e}_2\|$ and $z : 1$. We can apply our scaling factor to each respective axis before transforming the stretched edge back into the original global basis ($\tilde{\mathbf{v}}_{ab}^* = \mathbf{B}\mathbf{S}\tilde{\mathbf{v}}_{ab}$), where \mathbf{S} is a 3-by-3 scaling matrix with $\mathbf{S}_{11} = \|\mathbf{e}_1\|$, $\mathbf{S}_{22} = \|\mathbf{e}_2\|$, $\mathbf{S}_{33} = 1$ and remaining entries to be zero. This process is repeated for both points a and b using their respective eigenvectors ($\mathbf{e}_{1a}, \mathbf{e}_{2a}, \mathbf{e}_{3a}$ & $\mathbf{e}_{1b}, \mathbf{e}_{2b}, \mathbf{e}_{3b}$). The average length of the resulting two edges is the new edge length (as a discrete approximation to the integration over the edge).

4.3.3 Correspondence generation

We now describe our method to find correspondences, as per Algorithm 1. In addition to local feature matching, which only works for regions with limited deformation, we further identify correspondences using automatic correspondence inference in an as-consistent-as-possible formulation. The formulation works even under substantial non-isometric deformation.

4.3.3.1 Local feature matching In our implementation **SHOT** [194] signatures have been used to produce a candidate set of correspondences. **SHOT** computes a

pointwise descriptor based on extrinsic shape properties. The method relies upon computing repeatable 3D local reference frames that are invariant to some noise. The local space is then partitioned by an isotropic spherical grid. Similarly to the scale-invariant feature transform [124], for a given keypoint, first-order differential information at neighbouring vertices is grouped into bins to produce a set of local histograms, which form the signature. Due to the use of a fixed isotropic spatial grid, **SHOT** is sensitive to anisotropic non-isometric deformation.

To reiterate, we have chosen **SHOT** for feature matching as it has acceptable performance [192, 117]; the focus of this chapter is not to seek the best feature descriptor. **SHOT** signatures are computed at two scales for each vertex on X and Y [194]. Vertices between the two shapes are matched based on the similarity of their signatures.

As with most low-dimensional feature descriptors, locally symmetric areas may produce similar local signatures. This kind of incorrect correspondence may be identified by checking its consistency with other well matched feature points. For regions with isometric deformations, well established methods based on diffusion pruning [192] can be used. We will discuss this along with our extended approach for non-isometric deformation later in Section 4.3.4.

4.3.3.2 Correspondence identification based on automatic correspondence inference In regions with substantial non-isometric deformation, matching of local feature descriptors (*e.g.*, **SHOT**) often fails. Figure 4.8 is an example of a large deformation. Here, we propose to use reliable correspondences from isometrically deforming regions as landmarks to help identify new candidate correspondences in these regions. The landmarks are automatically selected and updated in each iteration of our surface registration framework. Given these landmarks, new candidate correspondences are proposed based on the consistency of their geodesic distances to the landmarks. There are three steps to this process, namely: problematic region identification, automatic landmark selection and correspondence matching.

Problematic region identification Initially, regions that contain substantial non-isometric deformation can be found by applying a pruning technique [192]. Regions with no (locally isometrically consistent) correspondences will be identified as problematic regions (Fig. 4.8). Each problematic region is then separately analysed. To reduce computation time, very small regions are disregarded since adequate correspondences will be identified later at the **ICP** stage.




Figure 4.8: Example problematic region between two shapes with no correspondences.

Automatic selection of landmark correspondences Correspondences returned from the pruning algorithm by Tam et al. [192] are generally reliable, especially those scoring high confidence values. In principle, we can randomly pick any of these correspondences as landmarks. However, landmarks that are far away from problematic regions are not very effective in estimating new candidate correspondences. Further, in practice, we observe that only a small number of landmarks are necessary to uniquely determine a new candidate correspondence point $\mathbf{y}_j \in Y$ for a given point $\mathbf{x}_i \in X$. We thus sample 20 landmarks that are closest to the boundary of a problematic region. This set of landmarks L helps to reduce the ambiguity and impact of slight errors in the landmark locations. Although a more sophisticated landmark selection procedure would be desirable, through experiments, we found that our method is insensitive to the number of landmarks, and produces consistent results when this parameter is varied within a reasonable range. We therefore fixed this parameter in our experiments.

Matching of correspondences For each point \mathbf{x}_i in a problematic region on surface X , the point $\mathbf{y}_j \in Y$ with the most consistent geodesic distances to landmarks is chosen as the corresponding point. Let $\mathbf{l}_k = (\mathbf{l}_k^X, \mathbf{l}_k^Y) \in L$ be the k^{th} landmark correspondence with $\mathbf{l}_k^X \in X$, $\mathbf{l}_k^Y \in Y$. We measure the consistency of a correspondence $(\mathbf{x}_i, \mathbf{y}_j)$ w.r.t. a landmark \mathbf{l}_k as

$$\text{cons}_k(\mathbf{x}_i, \mathbf{y}_j) = \min \left(\frac{d_g(\mathbf{x}_i, \mathbf{l}_k^X)}{d_g(\mathbf{y}_j, \mathbf{l}_k^Y)}, \frac{d_g(\mathbf{y}_j, \mathbf{l}_k^Y)}{d_g(\mathbf{x}_i, \mathbf{l}_k^X)} \right). \quad (4.4)$$

Similar to [82, 192], Eqn. 4.4 provides a normalised measure of consistency in the range of $[0, 1]$ where 0 is the worst and 1 is the best. $d_g(\cdot, \cdot)$ measures the geodesic distance between two vertices. However, unlike existing work that sets a threshold, which does not generally work for non-isometrically deforming regions, we choose the target corresponding vertex \mathbf{y}_j as the one that gives the overall best geodesic consistency. It works even for regions with non-isometric deformations:

$$L(x_i) = \arg \max_j \min_k \text{cons}_k(\mathbf{x}_i, \mathbf{y}_j). \quad (4.5)$$

The rationale is to choose correspondences that give good alignment (by maximising the minimum consistency) to all the landmarks, whilst reducing the effect of far away landmarks that are not discriminative enough for identifying a new good correspondence.

4.3.4 Correspondence pruning

Tam et al. [192] propose an efficient algorithm capable of pruning a set of candidate correspondences, preserving only those that are globally consistent. Poor correspondences are pruned based on their low consistency with good local correspondences. The consistency measure k_{ab} for a pair of correspondences $a = (\mathbf{s}_i, \mathbf{t}_u)$ and $b = (\mathbf{s}_j, \mathbf{t}_v)$ is defined as:

$$k_{ab} = \min \left(\frac{d_g(\mathbf{s}_i, \mathbf{s}_j)}{d_g(\mathbf{t}_u, \mathbf{t}_v)}, \frac{d_g(\mathbf{t}_u, \mathbf{t}_v)}{d_g(\mathbf{s}_i, \mathbf{s}_j)} \right), \quad (4.6)$$

where $\mathbf{s}_i, \mathbf{s}_j \in X$ and $\mathbf{t}_u, \mathbf{t}_v \in Y$ are vertices on surfaces X and Y respectively. In their work Huang et al. [82] and Tam et al. [192], a threshold c_0 (set to 0.7, as suggested by Huang et al. [82]) is introduced, and a pair of correspondences is considered acceptable if $k_{ab} \geq c_0$ in order to penalise non-isometric deformations. Tam et al. [192] extend [82] by considering only local isometry where end points \mathbf{s}_j and \mathbf{t}_v of a correspondence b must be within a specified geodesic distance to the end points \mathbf{s}_i and \mathbf{t}_u of a correspondence a . The global consistency of all reliable correspondences can be further inferred from local isometry via diffusion processing. This technique is shown to perform well in near-isometric cases.

However, in general, these techniques [82, 192] fail if the surfaces undergo non-isometric deformation, especially the large-scale deformations that are considered in this chapter. Our observation is that large-scale deformation often consists of non-isometric (esp. anisotropic) deformation. As these techniques model isotropic (near-)isometric deformation only, they fail to return any good correspondences essential for accurate **N-ICP** alignment. Here, we develop a local anisotropy metric

(Section 4.3.2) to estimate anisotropic deformation during registration, and apply Eqn. 4.6 to incorporate an anisotropic geodesic measure [119] $d_g^A(\cdot, \cdot)$. To our knowledge, this is the first effort to explicitly model anisotropic non-isometric deformation in a pruning technique.

4.3.5 N-ICP with extended ARAP regularisation

Our **N-ICP** implementation is inspired by Bouaziz and Pauly [27]. At each iteration of **N-ICP**, we work out the deformed position \mathbf{x}'_i for each vertex \mathbf{x}_i minimising an objective function involving both the data and regularisation terms, *i.e.*,

$$E = E_{\text{data}} + \lambda E_{\text{reg}}. \quad (4.7)$$

For the first 5 iterations, we only consider correspondences from pruning (C_{pruned} in Algorithm 1). This transforms the surface to an initial deformation that is more suitable. Let C be the set of pruned correspondences and $m = |C|$ be the number of correspondences input. Denote each correspondence as $c_k = (\mathbf{x}_{c_k}, \mathbf{y}_{c_k})$ where $\mathbf{x}_{c_k} \in X$ and $\mathbf{y}_{c_k} \in Y$ respectively. The data term we used is the standard formulation:

$$E_{\text{data}} = \sum_{k=1}^m \|\mathbf{x}'_{c_k} - \mathbf{y}_{c_k}\|_2^2. \quad (4.8)$$

After the 5th iteration, we discard prior correspondences, and use nearest neighbour correspondences to help guide the mesh into alignment. Our data term is now formulated as a combination of point-to-point and point-to-plane distances, where \mathbf{y}_i is the closest point to \mathbf{x}'_i on shape Y , and \mathbf{n}_i is the normal of point \mathbf{y}_i . w_1 and w_2 are the respective weights of point-to-point and point-to-plane terms.

$$E_{\text{data}} = w_1 \sum_{i=1}^n \|\mathbf{x}'_i - \mathbf{y}_i\|_2^2 + w_2 \sum_{i=1}^n \|\mathbf{n}_i (\mathbf{x}'_i - \mathbf{y}_i)\|_2^2 \quad (4.9)$$

For the regularisation term, traditional **N-ICP** uses a local **ARAP** energy in one-ring neighbourhoods [185]. We observe that this regularisation is not sufficient for registration of regions with large-scale deformations. For the purpose of shape deformation, Chen et al. [38] generalise this concept to work with r -ring neighbourhoods. It would provide greater control over how the deformation of a vertex affects its connected vertices. We propose to incorporate r -ring **ARAP** into our **N-ICP** registration pipeline to handle large-scale deformations. Let N_i^r be the r -ring neighbourhood for the vertex \mathbf{x}_i , and $n = |X|$ be the total number of vertices of X .

$$E_{\text{reg}} = \sum_{i=1}^n \sum_{j \in N_i^r} \|(\mathbf{x}'_j - \mathbf{x}'_i) - \mathbf{R}_i(\mathbf{x}_j - \mathbf{x}_i)\|_2^2 \quad (4.10)$$

The optimisation of E involves alternating optimisation of the deformed positions $X' = \{\mathbf{x}'_i\}$ for each vertex \mathbf{x}_i (while fixing \mathbf{R}_i), and the calculation of the optimal local rigid rotation \mathbf{R}_i (while fixing X'). Similar to [185, 38], the former (global step) is a least-squares problem and can be solved by solving a linear system, and the latter is optimised for each vertex and can be solved using Singular Value Decomposition [185, 38].

Note that on line 12, Algorithm 1, we pass X as the source shape rather than X' , this enables the consistent application of local regularisation and helps ensure better registration results.

4.4 Evaluation

To evaluate the performance of our proposed method, we benchmark it on a series of synthetic and real datasets. These datasets contain challenging registration scenarios, that include a range of minor and major anisotropic non-isometric deformations. We compare our technique against the notable **N-ICP** [27], two recent sparse non-rigid registration methods [210, 109], and a state-of-the-art functional map method [200].

4.4.1 Error measure

For convenience, the protocol of Kim et al. [95] is described here. For an estimated correspondence $(\mathbf{x}_i, \mathbf{y}_i) \in X \times Y$ and the respective ground-truth correspondence $(\mathbf{x}_i, \mathbf{g}_i) \in X \times Y$. The geodesic distance between the corresponding points on Y is $d_Y(\mathbf{y}_i, \mathbf{g}_i)$ (an example of geodesics is shown in Fig. 4.9). The area of shape Y is used to normalise the distance. The error of the estimated correspondence is measured as

$$\varepsilon(\mathbf{x}_i) = \frac{d_Y(\mathbf{y}_i, \mathbf{g}_i)}{\text{area}(Y)^{1/2}}. \quad (4.11)$$

Cumulative error curves are subsequently produced by counting the number of correspondences with an error $\varepsilon(\mathbf{x}_i)$ less than a given threshold of normalised geodesic distance ϵ , *i.e.*, $\varepsilon(\mathbf{x}_i) \leq \epsilon$.

[210, 109, 200] and the proposed method each rely on an initial set of correspondences. To ensure tests were fair, we computed an initial set of correspondences using **SHOT** [194]. Vestner et al. [200] was provided a dense set of feature descriptors; while, for [210, 109] and our method, we subsequently applied pruning [192] once for each pair of shapes, and used the same set of initial correspondences. To




Figure 4.9: Exemplar illustration of geodesic distance. For a point at the centre of the leopard’s face, the geodesic distance to any given point on the surface is indicated by the local colour value.

demonstrate the generalisability of each method, we do not perform further parameter tuning between tests, except where specified.

4.4.2 Implementation

A 352-dimension **SHOT** feature descriptor was used for [210, 200, 109] and our method. For the underlying **N-ICP** of our method, we used the following parameters: $w_1 = 10$, $w_2 = 2$, $\lambda = 1000$, iterations = 30, & r -ring = 2. The proposed PSF and correspondence inference methods were enabled in all tests, except where specified. For the **N-ICP** method [27], we kept most of the default settings, except the number of iterations, for which we found a value of 50 was sufficient. To run [200], we created a low resolution (5,000 vertices) copy of each mesh using *qs-lim* [70] and then applied *MeshFix* [11], providing both a high resolution and low resolution version of each shape during run-time. Shape pairs were rescaled to ensure their area was consistent. We enabled partial matching and set $\alpha = 10^{-7}$ and $t = [500, 323, 209, 135, 87, 56, 36, 23, 15, 10]$; initially solving for 10,000 correspondences, with a maximum problem size of 3,000 correspondences for all subsequent iterations.

4.4.3 Synthetic datasets

TOSCA high-resolution dataset TOSCA is a collection of 80 synthetic meshes of animals that range from mammals to mythological creatures. For each of the seven species, meshes of the same shape in different poses are provided. For the TOSCA dataset [32], we generated 69 pairs of shapes with the same class label (*e.g.*, cat to

cat, dog to dog, etc.). Each mesh was simplified to have 10,000 faces using *qslim* [70]. Models with flipped normals were excluded from tests. We disabled partial matching for [200] on this dataset, and kept other settings (see Section 4.4.2).

Fig. 4.10 shows the average registration accuracy of each method on the TOSCA dataset. These graphs show the proportion of vertices that have an error less than the value on the x -axis. The techniques corresponding to curves towards the top left of the graph are comparatively better. It should be noted that [27] failed to perform well on the TOSCA dataset, as well as the subsequent datasets. Since this was the simplest dataset, we excluded this method from further tests.




Figure 4.10: TOSCA high-res dataset results.

Fig. 4.11 shows the location of registration errors, a colourbar is included with each example to indicate the error value each colour represents. For every vertex, we measure the average registration error across a set of deformed models for both our proposed method and Yang et al. [210]. Fig. 4.12 shows the registration result when dealing with a large non-isometric intra-class deformation on the hind legs of the TOSCA dog.

Bouncing dataset Fig. 4.13 shows the results of our proposed method, compared with [210], on the Bouncing dataset [202]. We measure the fitting error between the deformed and target shapes with Hausdorff distance in MeshLab [40]. Note that [200] only produces correspondences rather than surface alignment, and cannot be used for this evaluation.

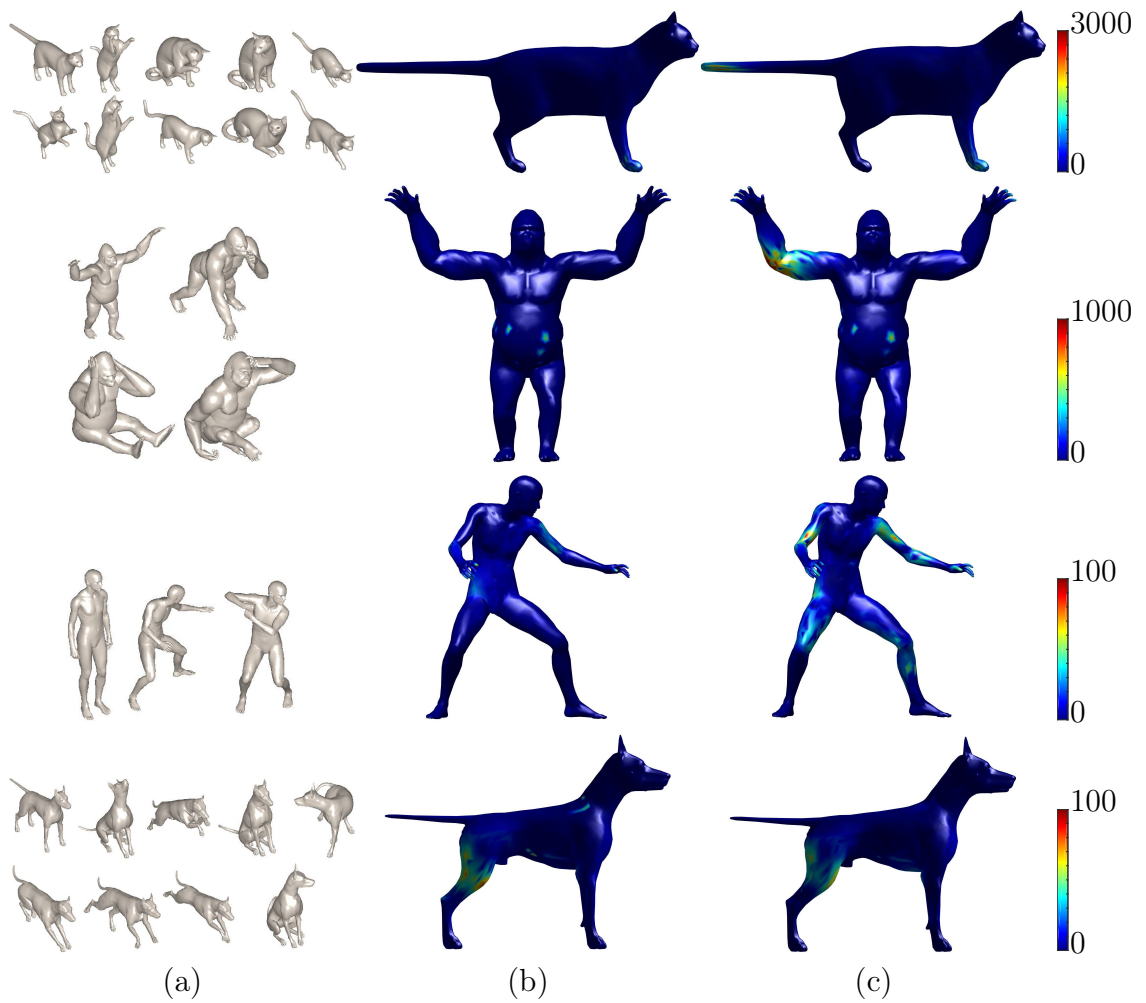


Figure 4.11: Results for the TOSCA cat, gorilla, David & dog sets, (a) source/target models used, (b) combined registration error of our method, (c) combined registration error of [210].




Figure 4.12: Registration results on TOSCA dog. (a) initial pose source (top) target pose (bottom), (b) our method (c) Yang et al. [210].




Figure 4.13: Deformation results on Bouncing dataset [202]. (a) initial pose source, (d) target pose (bottom), (b,e) our method (c,f) sparse non-rigid registration method [210].

SHREC’16 partial model dataset We further used the SHREC’16 partial model dataset [43]. This dataset is an important derivative work of the TOSCA dataset in which shapes have missing data (holes or cuts) that make non-rigid registration challenging. Fig. 4.14 illustrates the results of the Partial Functional Maps [159] and Random Forests [158] reported by Cosmo et al. [43] compared with the results of [210, 109, 200] and our proposed method. This demonstrates that our method is robust to holes and cuts, and clearly outperforms state-of-the-art methods, especially for the cases with holes.

To demonstrate the effectiveness of the individual ideas in our technique (namely, selection of landmarks, principal scaling factor, and r -ring **ARAP**), we ran a series of experiments with different configurations on the SHREC’16 partial model dataset. Results are illustrated in Fig. 4.15, which show that all the components in our method contribute to improved performance.

In summary, the results demonstrate that the proposed method consistently outperforms [210, 109, 200, 27]. In cases where the initial **SHOT** correspondences were poor, we found that the landmark-based correspondences and r -ring **ARAP** helped improve the alignment of large deformations.

4.4.4 Real datasets

Partial body scans Results from running our registration pipeline on a collection of partial body scans [4] are shown in Fig. 4.16. We also observed that missing data can cause our correspondence generation technique to produce more incorrect correspondences. In our tests, using two- or three-ring neighbourhoods helped alleviate this problem and improved registration results. The results presented in Fig. 4.16 were computed using three-ring neighbourhoods.

FAUST dataset We used the FAUST dataset [21] to help us objectively evaluate our method. FAUST is a collection of real scans captured using a 3D multi-stereo system. Ground-truth correspondences were acquired by covering each subject’s body with a series of stamps that were matched using texture-based registration. We compared the performance of our method, [210, 109] and [200].

Since none of the methods require any form of training, we used the training set for our evaluation. In our first experiment we produced an intra-subject test-set as follows: for each person in the dataset, we registered every mesh of that same individual to their first mesh, creating a set of 90 intra-person shape pairs. We then created our own inter-subject data by pairing each person in the dataset to another

person in the dataset in a single, randomly selected pose. This created a set of 90 randomly chosen pairs to evaluate each method with. We used exactly the same set of shape pairs for each method.

Registration results on the FAUST dataset are shown in Fig. 4.17, and some visual examples in Fig. 4.18. Our technique consistently shows lower mean geodesic errors compared against the state-of-the-art methods.

It should be noted that the curves in the respective figures only show the aggregated performance over the entire dataset. For simpler cases (shape pairs and regions), all the methods work well. When it comes to challenging regions (esp. joints), our technique performs significantly better than the curves show. The superior performance of our method is illustrated by typical examples in the paper.

4.5 Limitations

With the help of local scaling metric estimation and additional correspondences from landmarks, the method is more robust to large-scale, especially non-isometric deformations in surface registration than existing methods, as demonstrated by extensive comparisons. However, there are still some limitations. The proposed method still relies on a few correct initial correspondences to start the iterative process, and may perform poorly if those correspondences are largely wrong. Currently, a simple strategy is used to fix all the parameters. Some parameters, such as the number of landmarks and the r -ring **ARAP**, may work more effectively when more sophisticated, adaptive algorithms are used; this has been left as future work.

4.6 Conclusions

In this chapter a novel pipeline capable of registering two 3D shapes with large and non-isometric non-rigid deformations is presented. A technique to estimate local anisotropic non-isometric deformation and compute reliable correspondences in areas under large-scale deformation is developed. A new consistency measure and automatic landmark selection method to support non-isometric consistency are described. The **ARAP** regularisation constraint is then extended to deform larger surfaces in a more uniform manner. Experimental results demonstrate that this technique performs well in challenging scenarios. Comparative evaluation also highlights scenarios where state-of-the-art methods fail.

This technique performs well in non-isometrically deformed regions, and we believe that the technique can be further developed to handle more challenging cases (*e.g.*, heterogeneous objects). In addition, the current implementation only takes a source and a target shape as input. It is possible to generalise this method to exploit more example shapes to achieve more accurate registration results.

4.7 Summary

This chapter identified that the problem of non-isometric registration has been studied minimally. A non-rigid registration pipeline that addresses a fundamental problem in a novel manner was presented. Several key changes to an isometric registration pipeline were suggested to support anisotropic deformation. A detailed description of each component of the novel registration pipeline was presented. To evaluate the described approach, it was compared against state-of-the-art registration and shape correspondence techniques. The evaluation performed on this method is extended in several of the subsequent chapters (Chapters 5, 6 & 7).

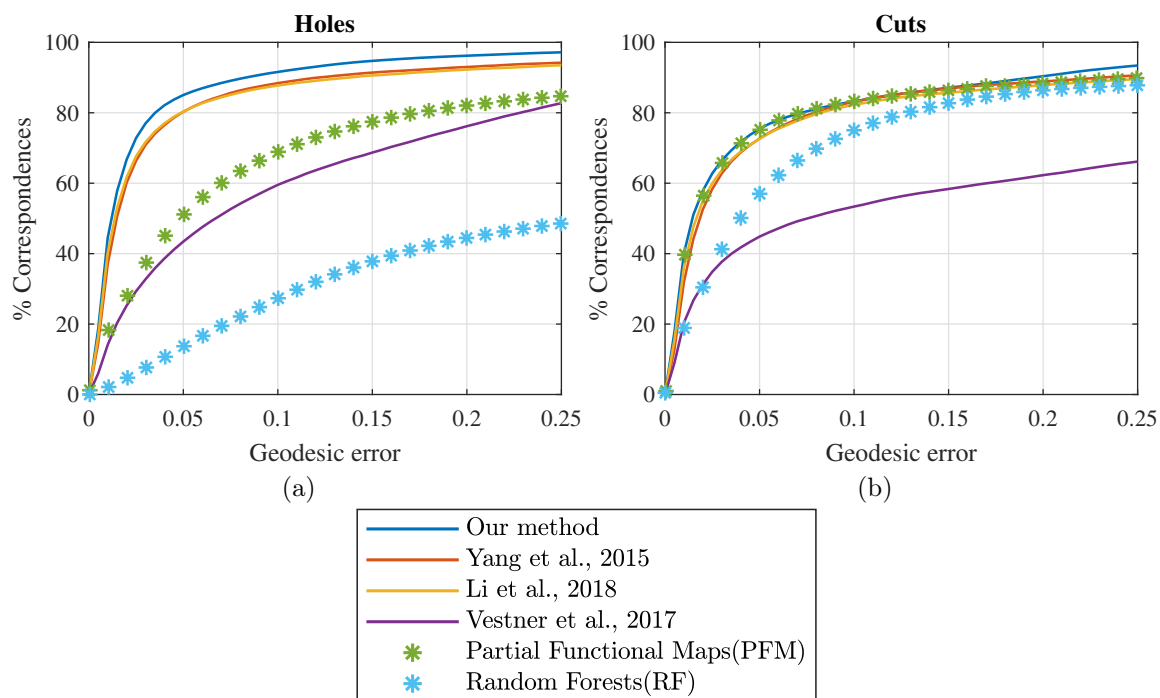


Figure 4.14: Results on SHREC'16 partial dataset compared with results (Partial Functional Maps & Random Forests) from the dense methods reported in [43]. The evaluation is split up into models containing (a) holes, and (b) cuts.

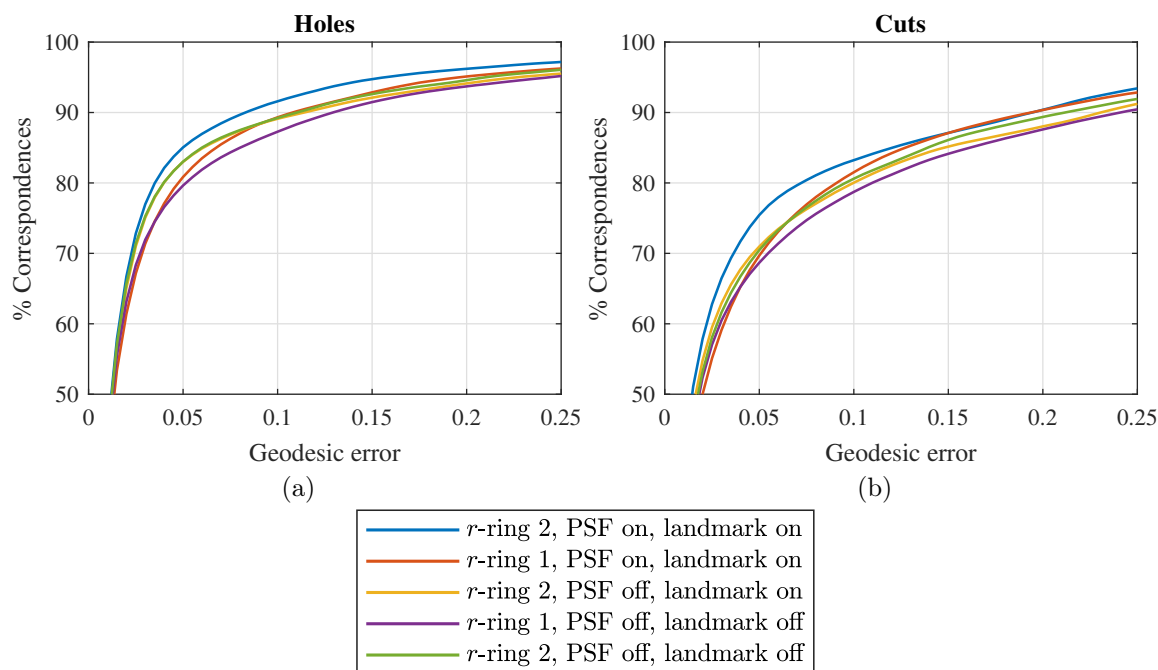


Figure 4.15: Results of our method in different configurations on SHREC'16 partial dataset. In configurations where PSF is off, standard geodesics are used, rather than anisotropic ones (skipping lines 7 & 8 of Algorithm 1). Where correspondence inference is off, we skip lines 10 & 11 of Algorithm 1. The evaluation is split up into models containing (a) holes, and (b) cuts.




Figure 4.16: Results of our proposed non-rigid registration technique on partial body scans [4]. (a,b,c) initial pose source X (blue), target Y (white), (d,e,f) overlapped result.

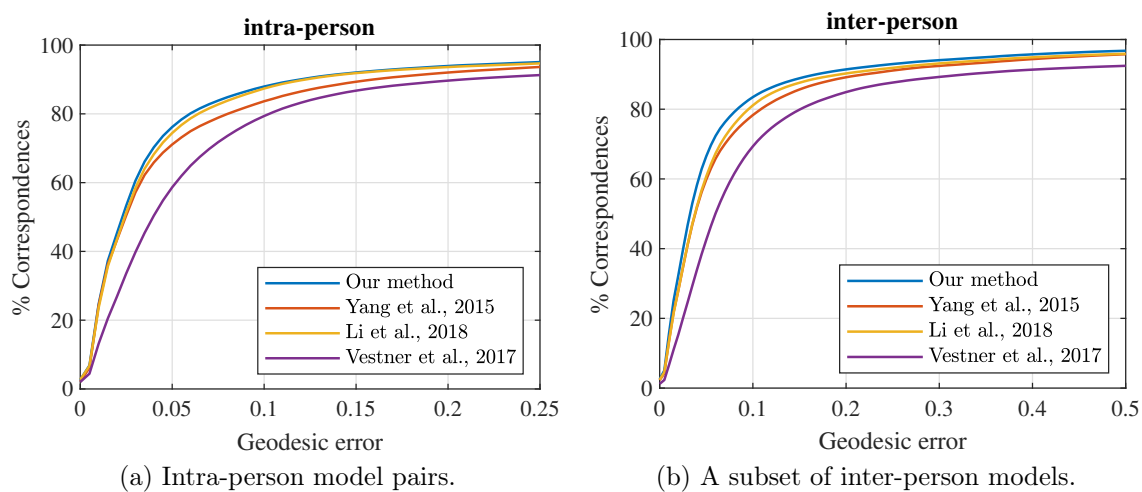


Figure 4.17: Registration results of pairs of models from the FAUST dataset.

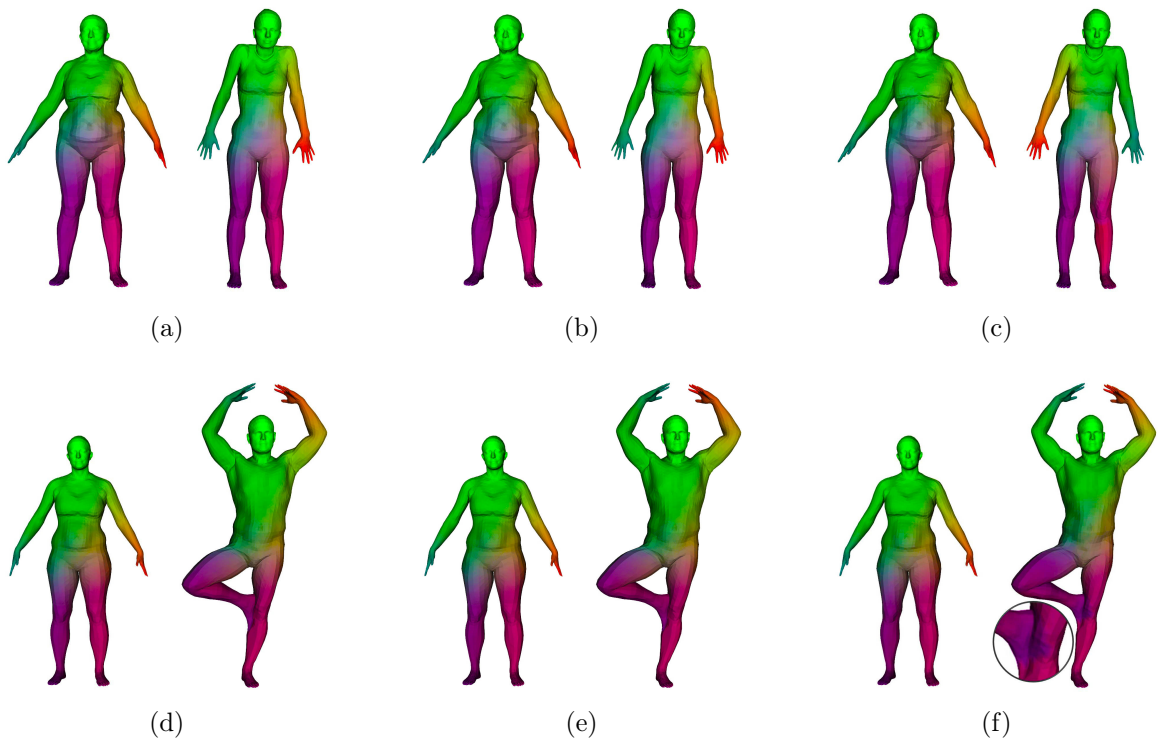


Figure 4.18: Screenshots of correspondence results of (a,d) our method, (b,e) Li et al. [109] and (c,f) Vestner et al. [200] on pairs of inter-person shapes from the FAUST dataset (left source, right target). Colours represent correspondence between shapes (*i.e.*, the same point on each shape should have the same colour). Note the reflection of correspondences of [200] in (c), and the colour bleed at the intersection between legs in (f). The mean geodesic error for each scan pair: (a) 0.0373, (b) 0.0419, (c) 0.2971, (d) 0.0639, (e) 0.0820, and (f) 0.0840.

Chapter 5

SHREC'19: Shape correspondence with isometric and non-isometric deformations




Figure 5.1: Examples of objects in the dataset.

Overview

This chapter presents a novel dataset for evaluating the performance of registration and shape correspondence methods on different types of non-rigid deformation. This chapter is organised as follows: Section 5.1 describes the work undertaken in this chapter, and clarifies the motivation behind producing a new dataset. Section 5.2 describes the dataset constructed and the approach taken to evaluate the results of submissions. Section 5.3 describes the configuration of the methods that are compared in this chapter. Section 5.4 evaluates the results obtained on each test set. Finally, Section 5.5 summarises this chapter with a conclusion of the findings of the work conducted.

5.1 Introduction

Chapters 5, 6 & 7 describe the work undertaken in a series of *SHape REtrieval Con-tests* (SHREC) events. The SHREC event focuses on shape retrieval challenges. One of the applications of surface registration is to facilitate 3D model retrieval; after the alignment of non-rigid shapes, it becomes easier to compare shapes since the correspondence between their elements is known. An important aspect of the SHREC event is the encouragement of collaboration. In this case, participants were invited to run their method on the dataset described in Section 5.2. They were then to submit their results along with a description of their approach. All further work (*e.g.*, conceptualisation, dataset production, evaluation, etc.), was completed as part of this thesis. The advantage of this approach is that it ensures the fair comparison between a wide range of methods. Such works provide a picture—albeit incomplete—of the practical performance of current state-of-the-art approaches.

Shape correspondence is necessary in common applications such as modelling [94], reconstruction [107] and tracking [140]. Many potential challenges can occur when estimating a correspondence between surfaces: deformations such as non-isometry (*i.e.*, stretching between shapes); ambiguities (*e.g.*, mapping features with little geometric detail); shape incompatibility [132] (*e.g.*, varying connectivity and vertex density); partial correspondence (*e.g.*, registering two incomplete scans); and semi-compatible shapes (*e.g.*, matching a human and a centaur). Most existing methods assume surface deformations to be either: piecewise rigid; (near-)isometric; and/or topologically consistent. As identified in the dataset taxonomy in Section 3.7.1, there are only a few public benchmark datasets for shape correspondence that challenge these assumptions about deformations [21, 43, 102, 7]. Previous contests by Cosmo et al. [43] and Lähler et al. [102] have used synthetic objects that produce deformations that are not realistic. [21, 7] do capture real-life objects, focusing on specific object categories (*i.e.*, human bodies and human faces), but neither benchmark suitably considers the large range of deformation that an object may undergo simultaneously. Instead of directly generating correspondences, non-rigid registration methods tend to produce a set of local transformations that deform one surface to align with the other. The proposed benchmark can also be used to evaluate such methods, by working out correspondences based on the deformed shape.

At the time of this work, there were no benchmark datasets that compartmentalised shape deformations into distinct sets. This has motivated the production of this new benchmark dataset of watertight shapes that is divided into distinct sets.




Figure 5.2: Illustrations of some of the challenges in the dataset.

Each set contains different types of deformation, ranging from piecewise rigid to non-isometric. The hierarchical structure of this dataset distinguishes it from previous works. An additional challenging set is included that introduces topological changes, which are ill-considered by other benchmarks. Structure and texture data is simultaneously captured; enabling ground-truth correspondences to be established by matching markers drawn on the physical objects.

5.2 Dataset

For this work a new dataset was produced from 3D scans of real-world objects, captured by ourselves using a high-precision 3D scanner (Artec3D Space Spider) designed for small objects. Each object exhibits one or more types of deformation. Surface deformations are classified into four distinct groups by level of complexity:

0. Articulating—piecewise rigid deformation
1. Bending—*isometric and near-isometric*
2. Stretching—*isotropic and anisotropic (e.g., Fig. 5.2a)*
3. Topologically changing—*heteromorphic (i.e., shapes of different topology. e.g., Fig. 5.2c)*

The dataset consists of wooden mannequins and wooden hands that are articulated. To produce other types of deformation, clothes have been created for the model from two materials. One material can bend but is resistant to stretching, and another can both bend and stretch. To induce greater non-isometry, modelling clay is used underneath the clothing of the mannequin model. Materials and objects

have been carefully selected to incrementally introduce these deformation types so that the limitations (w.r.t. deformation type) can be clearly identified. Because the dataset consists of real-world scans, it contains geometric inconsistencies and topological changes due to self-contacts. The real-scans also contain natural noise, varying triangulation of shapes and occluded geometry (*e.g.*, Fig. 5.2b). Some examples of models included in this dataset are shown in Fig. 5.1.

A total of 76 shape pairs were selected for the test sets (Table 5.1). Test-set 0 contains 14 pairs of articulating wooden hand objects that exclusively exhibit piecewise rigid deformation. Test-set 1 contains 26 pairs of models, comprising clothed mannequins and hands that deform (near-)isometrically. Test-set 2 contains 19 pairs of models; the pairings are between a thin clothed mannequin and a larger mannequin, ensuring significant non-isometry. Test-set 3 contains 17 carefully selected pairs that contain challenging geometric and topological changes.

Set name	No. of pairs	Model materials
Test-set 0	14	wooden hands
Test-set 1	26	clothed hands, clothed mannequins
Test-set 2	19	very stretched clothed mannequins
Test-set 3	17	all materials

Table 5.1: Test set structure.

Information about the dataset presented here, including how to access it, can be found in the Cardiff University data catalogue (doi: [10.17035/d.2019.0072003316](https://doi.org/10.17035/d.2019.0072003316)).

5.2.1 Ground-truth construction

To generate ground-truths, clearly drawn texture marks (*e.g.*, Fig. 5.3) were made on the surfaces of the objects used. Correspondences were initially determined automatically using the shape texture maps, and then manually corrected by multiple annotators to ensure ground-truths were accurate for this work (see examples in Fig. 5.6 of the obtained ground-truth).

A semi-automatic approach was devised to match corresponding marker points. First, a linear discriminator was manually trained to classify markers of different colours in the CIELAB colour space. An additional category to filter out points that did not belong to any marker was also used. For each marker, the texture patch was extracted. Then, the marker’s location, colour, and number of spokes were recorded. The processing steps used are shown in Fig. 5.4. By combining this



Figure 5.3: A photo of the real wooden hand used in the dataset after markers were drawn.

information with the information of neighbouring markers (as depicted in Fig. 5.5), a *query string* that uniquely describes a marker's position can be constructed. Due to discrepancies caused by non-rigid deformation, the order of the respective strings for a pair of scans may slightly differ. Therefore, candidate correspondences were produced by matching markers based on their *edit distance*, which was used as a measure of similarity. Edit distance is a metric used to quantify the similarity (or dissimilarity) of two ordered lists of nominal data (*i.e.*, symbols), which are referred to as strings. Given a query string and a target string, the edit distance is defined as the number of edit operations (insertion, deletion, or substitution) required to transform from the query string to the target string. When using edit distance, the sequences around the query point were treated as circular, so the minimal distance with any starting point was considered. Finally, diffusion pruning [192] was used to remove any geodesically inconsistent matches. Stride [186] provides further description and analysis of the approach.




Figure 5.4: An illustration of the processing pipeline applied after extracting (a) a local patch is projected to a 2D plane, (b) a linear discriminator is applied to classify pixels, (c) morphological operations (*e.g.*, dilating and pruning) are then applied, (d) the final skeleton is then used to count the number of endpoints.




Figure 5.5: An illustration of the construction of a query string using the points neighbouring a query point (circled with two rings). The query string encodes the distance (by order), colour and number of spokes.

5.2.2 Evaluation

Similarly to other shape correspondence benchmarks [43, 102], the correspondence quality of each method is evaluated using the evaluation procedure employed by Kim et al. [95]. The quality of shape correspondence has been evaluated automatically by measuring normalised geodesics between the ground-truth and predicted correspondence. Specifically, let $(\mathbf{x}_i, \mathbf{y}_i) \in X \times Y$ be a pair of corresponding points between surfaces X and Y , the normalised geodesic error ϵ_i between the predicted correspondence \mathbf{y}_i and the ground-truth position \mathbf{g}_i on surface Y is measured as:

$$\epsilon_i = \frac{d_Y(\mathbf{y}_i, \mathbf{g}_i)}{\text{area}(Y)^{1/2}}. \quad (5.1)$$

The following measurements are used to evaluate the performance of each method:

- An overall error measurement, for methods that complete all test sets.




Figure 5.6: Shape pairs from test-sets 0, 2 & 3 with ground-truth correspondences visualised. Matching colours between shape pairs represent corresponding points.

- Four plots (one for each test set) of cumulative geodesic error to demonstrate the performance of methods for individual types of deformation. This is also useful for participants that have not submitted results for all test sets.

5.3 Correspondence methods

This section presents the configuration of each approach used to find correspondences on one or more of the test sets. Seven methods were evaluated using the benchmark, namely: traditional N-ICP [27], anisotropic non-rigid registration [55], deep learning-based shape correspondence [74], non-isometric partial functional maps [200], non-rigid registration with reweighted sparsity [109], genetic optimisation-based (near-)isometric shape correspondence [172], and a commercial non-rigid registration tool [166]. The pipeline of each method evaluated is described in Section 3.5, a reference to each method is given in Table 5.2.

5.3.1 Traditional Non-Rigid ICP (N-ICP)

To provide an effective baseline to compare the performance of the recently developed approaches, a version of the well known **N-ICP** method by Bouaziz and Pauly [27] that extends the original rigid formulation of ICP [19] is used.

The following parameters were used $w_1 = 1$ (point-to-plane term), $w_2 = 1$ (point-to-point term), $w_3 = 1$ (global rigidity term), $w_4 = 1000$ (local rigidity term) and $iter = 100$, and shapes were kept at their original scale.

Reference	Section	Method
[27]	3.5.1	S. Bouaziz and M. Pauly. Dynamic 2D/3D registration for the Kinect. In <i>ACM SIGGRAPH 2013 Courses</i> , SIGGRAPH '13, pages 21:1–21:14. Association for Computing Machinery, July 2013. ISBN 978-1-4503-2339-0. doi: 10.1145/2504435.2504456
[55]	3.5.2	R. M. Dyke, Y.-K. Lai, P. L. Rosin, and G. K. L. Tam. Non-rigid registration under anisotropic deformations. <i>Computer Aided Geometric Design</i> , 71:142–156, May 2019. ISSN 0167-8396. doi: 10.1016/j.cagd.2019.04.014
[74]	3.5.3	T. Groueix, M. Fisher, V. G. Kim, B. C. Russell, and M. Aubry. 3D-CODED: 3D correspondences by deep deformation. In <i>European Conference on Computer Vision</i> , ECCV '18. Springer, September 2018. doi: 10.1007/978-3-030-01216-8_15
[200]	3.5.4	M. Vestner, Z. Löhner, A. Boyarski, O. Litany, R. Slossberg, T. Remez, E. Rodolà, A. M. Bronstein, M. M. Bronstein, R. Kimmel, and D. Cremers. Efficient deformable shape correspondence via kernel matching. In <i>International Conference on 3D Vision</i> , 3DV '17, pages 517–526. IEEE, October 2017. doi: 10.1109/3DV.2017.00065
[109]	3.5.5	K. Li, J. Yang, Y. K. Lai, and D. Guo. Robust non-rigid registration with reweighted position and transformation sparsity. <i>Transactions on Visualization and Computer Graphics</i> , 25(6):2255–2269, June 2018. doi: 10.1109/TVCG.2018.2832136
[172]	3.5.6	Y. Sahillioglu. A genetic isometric shape correspondence algorithm with adaptive sampling. <i>Transactions on Graphics</i> , 37(5):175:1–175:14, 2018. doi: 10.1145/3243593
[166]	3.5.7	Russian3DScanner. Wrap 3.3, April 2018. URL https://www.russian3dscanner.com/

Table 5.2: The methods evaluated by the dataset described in this chapter.

5.3.2 Non-Rigid Registration with Anisotropic Estimation

The approach by Dyke et al. [55]—described in the previous chapter—is used. All parameters remain unchanged. Sparse correspondences computed by [192] were used to initialise the method.

5.3.3 3D-CODED

The hyperparameters for the deep learning matching approach by Groueix et al. [74] were unchanged for all the tests. The code used is available online from the authors of 3D-CODED [73] and the pre-trained network called: “sup_human_invY_network_last.pth” was used. An example output of the deformed template using this method is shown in Fig. 5.7.




Figure 5.7: This figure shows an example output of the deformed template using [74].

5.3.4 Efficient Deformable Shape Correspondence via Kernel Matching

The method proposed by Vestner et al. [200] was run by the authors using the publicly available code [103]. Two sets of results were submitted for this method with different parameters. The diffusion parameters for heat kernels remained the same throughout the benchmark. After normalising the shapes to have unit surface area, the diffusion parameters are set to

$$\underbrace{(\underbrace{\log_{10}(500), \dots, \log_{10}(500)}_{3x}, \dots, \log_{10}(10), \dots, \log_{10}(10)})_{10, \text{ logarithmic sampling}}}_{3x}. \quad (5.2)$$

Parameters for SHOT results As pointwise descriptors **SHOT** [194] is used as described in the paper.

Parameters for SHOT & HKS results **SHOT** and **HKS** [189] are used as pointwise descriptors.

5.3.5 Reweighted Position and Transformation Sparsities

For the method by Li et al. [109], the diffusion pruning [192] method is used to obtain initial correspondences based on matching of local **SHOT** features [194], and further correspondences are introduced during iterative optimisation based on closest points, similar to the standard N-ICP framework. All parameters were optimised by the original authors.

5.3.6 Genetic Isometric Shape Correspondence

For Sahillioğlu [172], two sets of results were submitted for this method with different levels of sparsity, one set that has relatively sparse correspondences (≈ 100 per shape pair) and one set of extremely sparse correspondences (6 per shape pair). A random result from each deformation type is shown in Fig. 5.8.




Figure 5.8: Extremely sparse (top) and sparse (bottom) correspondences produced by Sahillioğlu [172] on some pairs.

5.3.7 RSDS Wrap 3.3

For consistency, the diffusion pruning [192] method is used to generate an initial set of correspondences for the initialisation of *Wrap 3* [166]. The method runs iteratively where the density of control nodes is increased per iteration. This leads to an approximation of the target shape with increasing accuracy per iteration.

5.4 Results

Quantitative results of the methods described in Section 5.3 are presented here. Figs. 5.9a to 5.9d show the geodesic error of correspondences generated by each method on each test set. Fig. 5.10 shows the combined results for methods that registered all test sets. Table 5.3 reports the percentage (and where appropriate number) of correspondences returned by each method. Table 5.4 shows the *area under the curve* (AUC) of each method on each test set.

Evaluation This work uses the same error measure described in Section 4.4.1 in Chapter 4.

Sparsity of correspondence results Most methods evaluated here compute a reasonably dense set of matches. The number of correspondences returned by each method is reported in Table 5.3, 100% indicates that a correspondence was found for all vertices of X . As discussed in Section 5.3, Sahillioğlu [172] submitted two sets of results, the first set consisting of an average number of 98.3 (to one decimal place) sparse correspondences, the other containing 6 correspondences. Methods [109, 55, 27, 166] produce a deformed source mesh towards the target mesh. The following strategy is used to calculate the correspondences: For each vertex on the deformed mesh, the point is projected onto a triangle on the target mesh surface, and the barycentric coordinates are recorded. Correspondences that have a projection distance larger than the average mesh edge length are rejected, as this indicates regions where surfaces are not accurately aligned. Before rejecting correspondences, it was noted that the overall results of “all test sets” for Li et al. [109] & Dyke et al. [55] were comparable.

Test-set 0 This set contains only articulated deformations. This is the most simple type of deformation that is investigated. Thus most methods are expected to perform well on this test set. However, performance was discovered to vary across many methods, especially when compared to the other test sets. Inspecting the shapes

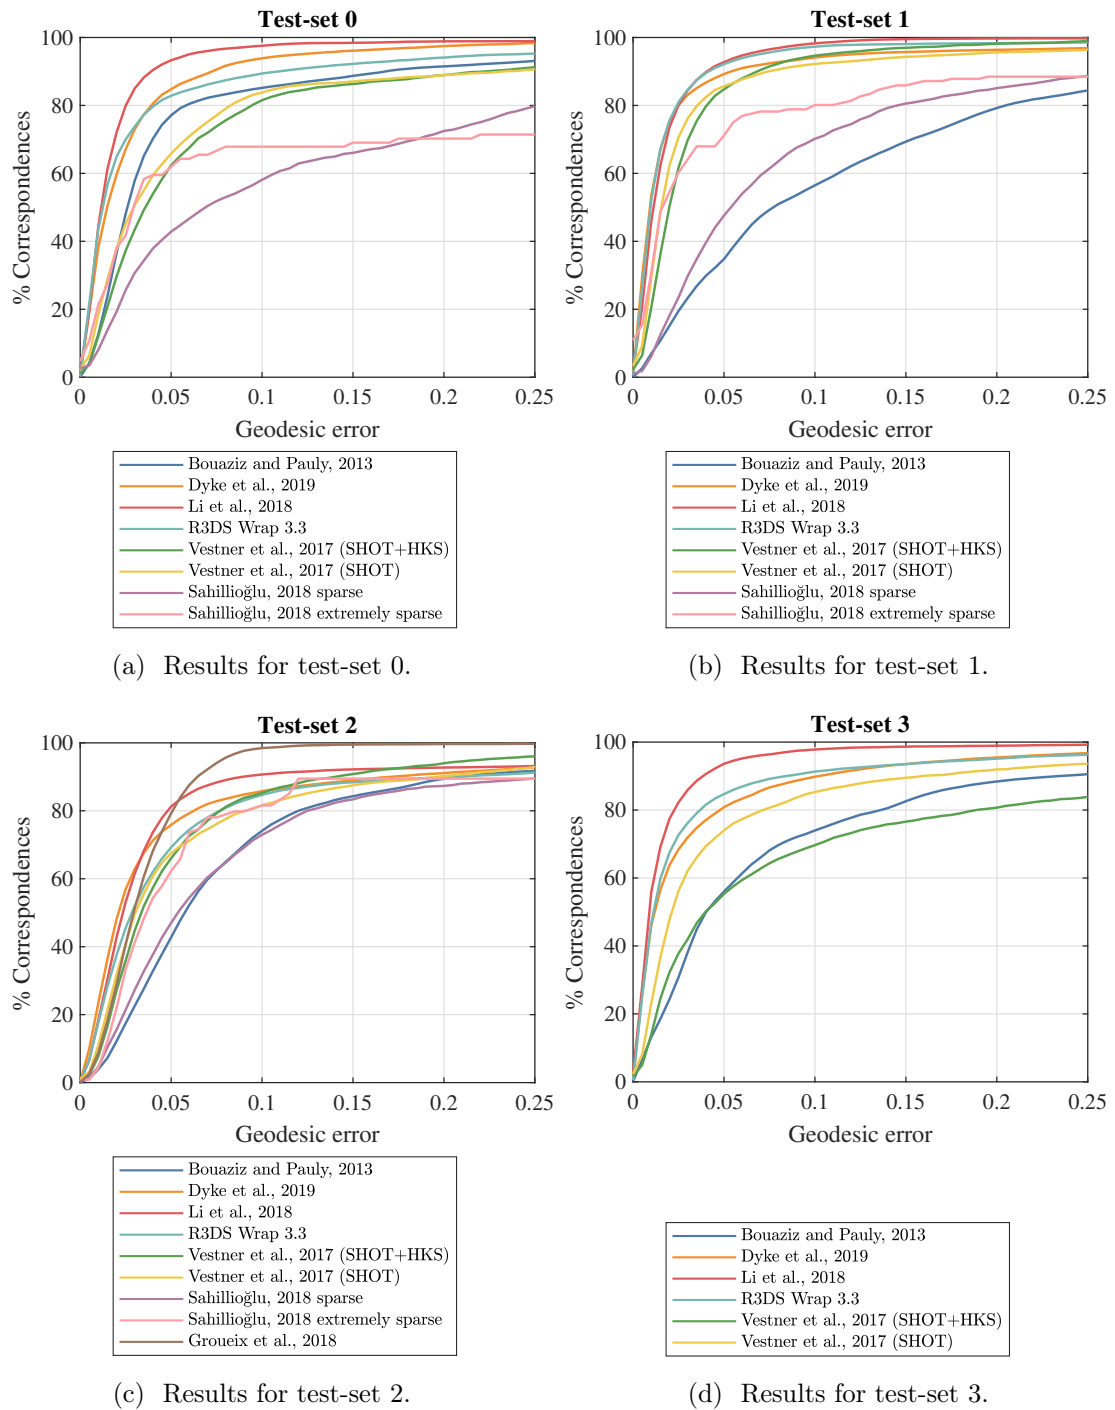


Figure 5.9: Results of all methods on each test set.

Method	Correspondences
Dyke et al. [55]	99.02%
Li et al. [109]	75.96%
Bouaziz and Pauly [27]	49.64%
Russian3DScanner [166]	93.99%
Vestner et al. [200] (SHOT+HKS)	92.31%
Vestner et al. [200] (SHOT)	92.39%
Groueix et al. [74]	100%
Sahillioğlu [172] sparse	98.34
Sahillioğlu [172] extremely sparse	6

Table 5.3: Sparsity of correspondence results given as either a percentage of the number of vertices comprising shape X or the number of correspondences, where appropriate.

in the dataset reveals that the surfaces are primarily comprised of smooth surfaces that lack high frequency geometric details. For example finger regions (cylinder-like surfaces) are symmetrically ambiguous (see Fig. 5.11). This may affect the initial correspondences leading to a higher error rate. [109] performs well because the large smooth surfaces fit the sparsity assumption.

Test-set 1 This test set contains shape pairs that bend either isometrically or near-isometrically. In it, the largest difference was observed between the best performing method (Li et al. [109]) and the worst performing method (baseline N-ICP [27]). Shape pairs also have large-scale deformations, which typical **N-ICP** methods ([27]) cannot handle as **N-ICP** requires two shapes to have a good initial alignment to ensure optimal registration.

Test-set 2 Groueix et al. [74] was observed to have the fastest convergence to 100%. It should be noted that this test set contains only non-isometric human-shaped models. [74] demonstrates how the use of a pre-trained network from some datasets may be generalised for other datasets. With respect to the other methods, observe that **SHOT**-based approaches suffer significantly, when compared with the results in test-set 0 and test-set 1. This is likely caused by the local non-isometry. As **SHOT** signatures are not well defined for such non-isometric surfaces, the degradation in performance is reasonable.

For Sahillioğlu [172], the performance degrades on shape pairs of mannequins that possess bilateral symmetry (test-sets 1 and 2). Due to self-occlusions during scanning, unnatural connections between fingers of some hand models are present, causing some

pairwise geodesics to be inconsistent in test-set 0. As a purely geodesic-based method, its performance is also affected by the unnatural shortcuts present in hand models due to occlusions during capture. It should be noted that the performance of [172] would have improved significantly if such problematic pairs (symmetric flips and shortcuts) were discarded.

Test-set 3 Vestner et al. [200] achieve notably worse results through the combination of **SHOT** and **HKS** when compared with using **SHOT** alone. This is contrary to the results of test-set 2, where using **SHOT** and **HKS** combined performs better. This demonstrates the instability of **HKS** under topological change. Topological changes present new challenges, and are likely to be beyond the assumptions of most methods. Therefore, some methods did not participate in this test. However, for those that participated in this test set, the overall **AUC** appears comparable to other test sets. This is probably because, apart from topological changes, this test set tends to have less distortion.

5.5 Conclusions

In this chapter a novel dataset of real-world scanned objects that cover a large variety of deformation types has been presented. Investigations have found that changes in topology is a challenging problem for some methods.

Machine learning-based approaches prove to be more capable of handling non-isometric deformations. However, they often require a high training cost, and may not generalise to arbitrary data. Recent advances in non-rigid registration techniques that explicitly model non-isometric deformation generally perform well in many scenarios. Though they do not perform as well as deep learning techniques in non-isometric deformation, they do not require training, are generically applicable to unseen datasets, and are less susceptible to topological changes. There is also a need to develop more reliable features for point-based correspondence on non-rigid surfaces.

To summarise, experimental results suggest that developing correspondence techniques that are generic and reliable to any kind of seen/unseen deformation and surface, whilst handling noise and topological changes, is still an on-going challenge. No single technique is perfect, but the results also indicate an interesting direction: combining the individual advantages of sophisticated deep learning models and the advantages of generic non-rigid non-isometric registration techniques may lead to a

Table 5.4: The resulting AUC for each method on each test set. Rows are in descending order of score over all test sets.

Method	test-set 0	test-set 1	test-set 2	test-set 3	all test sets
Li et al. [109]	0.9198	0.9259	0.8239	0.9286	0.8992
Dyke et al. [55]	0.8784	0.8986	0.8011	0.8583	0.8615
RSDS Wrap 3	0.8532	0.9199	0.7725	0.8698	0.8595
Vestner et al. [200] (SHOT)	0.7601	0.8651	0.7566	0.7992	0.8039
Vestner et al. [200] (SHOT+HKS)	0.7391	0.8710	0.7763	0.6638	0.7767
Bouaziz and Pauly [27]	0.7901	0.5720	0.6750	0.7033	0.6673
Sahilliođlu [172] sparse	0.5654	0.6595	0.6743	–	–
Sahilliođlu [172] extremely sparse	0.6300	0.7763	0.7459	–	–
Groueix et al. [74]	–	–	0.8605	–	–

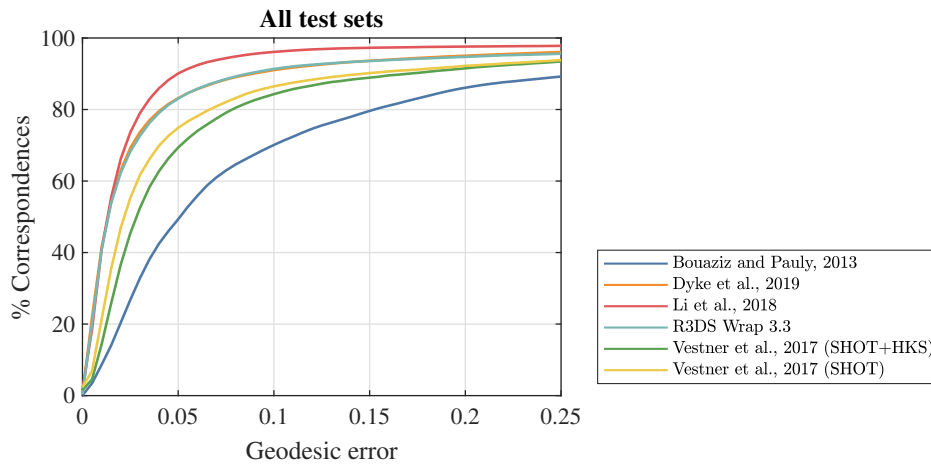


Figure 5.10: Results for all test sets.

more useful and generic correspondence technique that performs well in most scenarios, and would practically be applicable in downstream applications.

Overall, R3DS Wrap 3 performed similarly to [55]. It is interesting to see how well a commercial solution compares to the state-of-the-art methods.

Through this work some challenges in the fair evaluation of the performance of shape correspondence methods have been discovered. Taking intrinsic symmetries of shapes into account and reporting details—such as the sparsity of correspondences estimated—need further investigation.

Further exploration of the robustness of shape correspondence methods on partial real scans would contribute valuable information to the field, and this dataset could be augmented to provide such challenges in the future.

5.6 Summary

This chapter identified deficiencies in current benchmarks. Critically, a lack of tiers in benchmarks can preclude valuable information from being extracted. With this in mind, a real dataset containing a range of deformations (*i.e.*, articulated, isometric, non-isometric & topological change) and an acquisition process for ground-truths was proposed. Numerous registration and shape correspondence methods were benchmarked on the novel dataset. The results of all methods were then evaluated, focusing on the performance in each test set. In the two subsequent chapters (Chapters 6 & 7), further evaluation is performed using novel datasets.




Figure 5.11: Wooden hand object from the dataset coloured by the surface normals to illustrate the lack of high frequency geometric detail on the shape's surface.

Chapter 6

SHREC’20: Shape correspondence with physically-based deformations

Overview

This chapter presents a novel dataset that comprises of extreme non-isometric deformations of a single real shape. The problem that this dataset seeks to address is introduced and the key contributions are clarified in Section 6.1. The dataset and the acquisition method used are described in Section 6.2. Additionally, a brief investigation into the problem of correspondence initialisation is undertaken in Section 6.2.1, the results of which are discussed in Section 6.4.1. Participants were asked to estimate the correspondence between each partial scan in the dataset and one watertight scan of the rabbit. Section 6.3 details the configuration of the correspondence methods evaluated using this dataset. The results submitted by participants are presented and evaluated in Section 6.4.2. Finally, Section 6.5 contains concluding thoughts from the outputs of this work.

6.1 Introduction

This chapter presents work conducted in a subsequent SHREC event to Chapter 5. This work was conducted in parallel with the work described in Chapter 7. Both works consider different aspects of non-isometric deformation—the distinct contribution of each work is established in the introduction of the respective chapters. Simply put, the non-isometric deformations considered in Chapter 5 are somewhat limited. The work in this chapter identifies both complex poses and various internal materials, which are overlooked in existing literature. While the work in Chapter 7 focuses on a

constrained set of non-isometric deformations on quadrupeds, by exploiting existing data sources and using sparse annotations by experts for benchmarking.

Many recent registration and shape correspondence works have focused on addressing the problem of non-isometric deformations [201, 200, 67, 55, 88, 66, 60]. However, the resources available to evaluate the performance of such methods are insufficient. Previous SHREC events have produced insightful benchmark datasets. Many use synthetic objects [30, 102, 159, 160]. While these synthetic datasets may have certain advantages, such as easily established ground-truths, they also produce deformations that are not necessarily realistic.

Furthermore, current evaluative datasets only consider a few aspects of shape deformation, such as missing data and non-isometry. By providing datasets that focus on other problems in surface registration and shape correspondence, state-of-the-art techniques can be developed to address new challenges. Although the dataset described in Chapter 5 makes significant progress by developing a dataset that considers a range of deformations, there remain other important aspects of method performance that should be analysed. Due to the use of human-based models, the range of deformations were constrained. The dataset proposed here does not use a skeletal model, therefore increasing the challenge of discerning the deformation induced. This dataset also considers the use of granular internal materials that affect both the deformation exhibited and the surface of the shapes.

This work seeks to address a novel aspect of handling problems such as non-isometry and missing data by designing a hierarchical, incrementally challenging dataset. We focus on inducing non-rigid deformations on a real-world object and use texture markers to establish ground-truths. Various internal materials and drastic poses of the shape were selected, causing complex deformations, and varying degrees of local protrusions and indentations on the surface. A subset of deformations were selected—twist, indent, inflate & stretch—from those identified by Schmidt et al. [179], whose work focuses on how humans interpret deformation of an unknown object. No previous works provide such a range of complex deformations. Also, to the best of the author’s knowledge, this is the only correspondence benchmark that considers the internal materials of the given shapes.

6.2 Dataset

For this work, 3D scans of a real-world object were captured using a high-precision 3D scanner (Artec3D Space Spider). The scanner is able to match the actual position

of an observed location at up to 0.05mm accuracy. Each scan exhibits one or more types of deformation. Scans were classified into four distinct groups by the type of deformation primarily exhibited by a given pose: *twist*, *indent*, *inflate*, & *stretch*. The challenge of the dataset was further increased by varying the internal properties of the object by filling it with different materials: *couscous*, *risotto*, & *chickpea*. Using different internal materials changed the local appearance of the shape’s surface, as well as its deformation properties. For each pose, the object was filled, and then scanned, with three different internal materials. This caused incremental changes to the surface and overall deformation exhibited (see Fig. 6.1). In Fig. 6.2, further variations such as the number of creases on the surface may be observed as the internal material changes.




Figure 6.1: Illustrations of the surfaces of meshes with different internal materials coloured by the surface normals. These materials help induce varying magnitudes and frequencies of protrusions and indentations on the surface.

The dataset consists of a stuffed soft toy rabbit made out of a stretchable jersey material with no internal structure or skeleton that could restrict its movement, see Fig. 6.3. The rabbit had 590 coloured markers drawn on the surface, which allowed numerous accurate ground-truth correspondences to be established, see Fig. 6.4. Note that our purpose is to investigate how different types of physically-based deformations affect non-rigid shape correspondence, so a carefully chosen object with different material fillings is sufficient and makes data capture and analysis more manageable.

The poses and materials used were carefully selected to incrementally vary the deformation challenge so that algorithm deficiencies—with respect to these varying properties—may be identified. Some examples of challenging problems are shown in Fig. 6.5.

A key point about the proposed dataset is that the exaggerated nature of the deformations, such as twisting, often contradicts the underlying assumptions of state-

of-the-art approaches. Due to changes of the internal material and surface creases, geometry is often locally non-isometric, which is problematic for many shape descriptors. The model has intrinsic symmetries and, with the exception of the target scan, all scans are partial.

Information about the dataset presented here, including how to access it, can be found in the Cardiff University data catalogue (doi: [10.17035/d.2020.0112374975](https://doi.org/10.17035/d.2020.0112374975)).

6.2.1 Initial correspondences

Most correspondence methods require an initial set of sparse or dense correspondences that are subsequently refined. Due to the challenging variations in both the local and global appearance, three sets of initial correspondences using different methods were produced to find an optimal approach. The following methods were used to acquire the initial sets of correspondences: diffusion pruning [192], region-level correspondence [100], and N-ICP [27]. All methods took a similar amount of time to compute. Where possible these initial correspondences have been used for fairer comparisons.

The diffusion pruning method by Tam et al. [192] was used to produce a set of globally consistent correspondences from an initial set of point-wise descriptors. Due to the significant deformation, local geometry is often non-isometric and thus local descriptors that rely on near-isometries performed poorly. **SHOT** [194] signatures were computed with 10 bins at two scales, which covered 2% & 5% of the surface area square rooted. Due to limitations of memory and computation time, diffusion pruning was performed with the default parameters, except $K = 5$ and $d = 25\%$.

The region-level correspondence method by Kleiman and Ovsjanikov [100] was used to produce a set of correspondences between segmented shapes. This method is capable of finding region-based matches using a shape-graph that describes the connectivity of consistently segmented shapes. The method was run using default parameters, with the exception of `numComponentsRange = {10, 9, 8, 7}`. A dense point-to-point mapping was subsequently recovered using nearest neighbours of a functional map from a state-of-the-art approach [142].

The non-rigid registration method by Bouaziz and Pauly [27] was used to register the shapes together. Point-to-point correspondences were computed using nearest neighbour between the vertices of the two surfaces. The method's parameters were set to $w1 = 1$ (point-to-plane term), $w2 = 1$ (point-to-point term), $w3 = 1$ (global rigidity term), $w4 = 1000$ (local rigidity term) and `iter = 100`.

6.3 Correspondence methods

This section presents the methods that are examined in this work. Eight methods were evaluated using the benchmark; namely: non-rigid registration with re-weighted sparsities [109], non-rigid partial functional maps [159], anisotropic non-rigid registration [55], non-rigid partial functional maps [200], precise recovery of non-isometric functional maps [67], continuous and orientation-preserving functional maps [155], data-driven non-rigid registration *FARM+*, and a commercial non-rigid registration tool *R3DS Wrap 3.4*.

To reduce duplication, the introductory description of each method evaluated is presented once in Section 3.5, a reference to each method evaluated here is given in Table 6.1. Relevant configuration details that are specific to the experiments undertaken here are still included.

6.3.1 Robust Non-Rigid Registration with Reweighted Position and Transformation Sparsity

For the method proposed by Li et al. [109], the set of pruned and **N-ICP** correspondences were provided, and was run using parameters optimised by the authors.

6.3.2 Partial Functional Correspondence

The functional mapping framework by Rodolà et al. [159] that seeks addresses the correspondence problem between partial isometric shapes was also evaluated.

For the experiments, shapes were re-scaled—as described in the method’s code. All parameters remained default, with the exception of `n_eigen = 100`. **SHOT** was used to compute dense point-wise descriptors.

6.3.3 Non-rigid Registration under Anisotropic Deformations

The proposed method by Dyke et al. [55] (described in Chapter 4) is fully automatic. However, for comparative purposes, our pre-computed correspondences are used for initialisation. All other parameters remain as described in the original paper.

6.3.4 Efficient Deformable Shape Correspondence via Kernel Matching

Based on the discussions and experimental results found in the original paper by Vestner et al. [200], the parameters α and t were set to 10^{-7} and

CHAPTER 6. SHREC'20: SHAPE CORRESPONDENCE WITH
PHYSICALLY-BASED DEFORMATIONS

Reference	Section	Method
[27]	3.5.1	S. Bouaziz and M. Pauly. Dynamic 2D/3D registration for the Kinect. In <i>ACM SIGGRAPH 2013 Courses</i> , SIGGRAPH '13, pages 21:1–21:14. Association for Computing Machinery, July 2013. ISBN 978-1-4503-2339-0. doi: 10.1145/2504435.2504456
[109]	3.5.5	K. Li, J. Yang, Y. K. Lai, and D. Guo. Robust non-rigid registration with reweighted position and transformation sparsity. <i>Transactions on Visualization and Computer Graphics</i> , 25(6):2255–2269, June 2018. doi: 10.1109/TVCG.2018.2832136
[159]	3.5.8	E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. <i>Computer Graphics Forum</i> , 36(1):222–236, January 2017. ISSN 0167-7055. doi: 10.1111/cgf.12797
[55]	3.5.2	R. M. Dyke, Y.-K. Lai, P. L. Rosin, and G. K. L. Tam. Non-rigid registration under anisotropic deformations. <i>Computer Aided Geometric Design</i> , 71:142–156, May 2019. ISSN 0167-8396. doi: 10.1016/j.cagd.2019.04.014
[67]	3.5.9	D. Ezuz and M. Ben-Chen. Deblurring and denoising of maps between shapes. <i>Computer Graphics Forum</i> , 36(5):165–174, August 2017. ISSN 0167-7055. doi: 10.1111/cgf.13254
[155]	3.5.10	J. Ren, A. Poulenard, P. Wonka, and M. Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. <i>Transactions on Graphics</i> , 37:248:1–248:16, December 2018. doi: 10.1145/3272127.3275040
[126]	3.5.11	R. Marin, S. Melzi, E. Rodolà, and U. Castellani. FARM: Functional automatic registration method for 3D human bodies. <i>Computer Graphics Forum</i> , 39(1):160–173, February 2020. doi: 10.1111/cgf.13751
[167]	3.5.7	Russian3DScanner. Wrap 3.4, June 2019. URL https://www.russian3dscanner.com/

Table 6.1: The methods evaluated by the dataset described in this chapter.

{500, 323, 209, 135, 88, 57, 37, 24, 15, 10} respectively. Only **SHOT** is used to compute point-wise descriptors, as **HKS** is sensitive to changes in topology. Optimal results were achieved by scaling input shapes with respect to the unit area of the source shape. Since all shapes were captured using the same device, the original scale

was the same.

6.3.5 Deblurring and Denoising of Maps between Shapes

The functional map method by Ezuz and Ben-Chen [67] was run in three configurations, using each respective set of pre-computed initial correspondences. Due to memory limitations experienced when computing the wave kernel map, all configurations used a subset of 200 correspondences—selected using *farthest point sampling* (FPS) based on geodesic distances. The following parameters were modified: $k_1 = 120$ & $k_2 = 120$.

6.3.6 Continuous and Orientation-preserving Correspondences via Functional Maps

The method by Ren et al. [155] was run both with and without our sets of initial correspondences. Similarly to [67], due to limitations of memory when computing the wave kernel map, 200 correspondences were selected using geodesic-based **FPS**. Parameters were configured to the default settings used in the paper, with `num_iters` = 10. The parameters for computing **WKS** were `numTimes` = 100 and `skipSize` = 10.

6.3.7 FARM+

A variant of [126], referred to here as *FARM+*, is used. Since no deformable model is provided, *scan00* is animated using Mixamo [2] and some deformation basis is defined to inflate or shrink the template along the direction of the surface normals. To compute landmarks, the minimum and maximum of the first Laplacian eigenfunctions are used, and six landmarks are classified over extremal points of the rabbit’s ears, arms, and legs. Another main change of the original methods by Marin et al. [125] and Marin et al. [126] is that no local correspondence is used. All parameters were left unchanged from the original method, which were tuned for the specific domain of human bodies.

6.3.8 R3DS Wrap 3.4

At the time of this work, *Wrap 3.4* [167] was the latest available version of the R3DS commercial software utility. The publisher’s documentation does not indicate any notable modification to the implementation described in Section 3.5.7. *Wrap 3* was used to register the scans both with and without our initial pruned correspondences.

Better results were achieved by setting the final weight of control points (the initial correspondences) to zero, all other settings remained as per default.

6.4 Results

In this section we discuss the performance achieved by correspondence methods on our dataset. We examine how different initialisation techniques impact the performance of a selection of methods. Then the effects of variations in pose and internal material are scrutinised. It should be noted that the density of the correspondences acquired is not wholly reported because all methods report a dense mapping of $\approx 100\%$ in the majority of cases.

Evaluation This work uses the same error measure described in Section 4.4.1 in Chapter 4.

6.4.1 Initial correspondences

In Fig. 6.6 the performance of the initial correspondence methods described in Section 6.2.1 is reported. The methods are abbreviated as follows: *N-ICP* [27] refers to [27], *Pruning* [192] refers to the diffusion pruning method by Tam et al. [192], and *SEG* [100] refers to the region-level correspondence method by Kleiman and Ovsjanikov [100]. N-ICP [27] performed the best. This may be due to the shapes having a reasonable initial rigid alignment, which is necessary when using N-ICP. All deformations are largely locally non-isometric; this causes the point-wise descriptors used to become unreliable.

These initial correspondences were subsequently used to initialise the relevant shape correspondence and registration methods. In Fig. 6.7, a comparison of the performance of a selection of methods using different initial correspondences is reported. The results show that using the initial correspondences of N-ICP [27] greatly reduces the resulting error of the subsequent methods.

Table 6.2 summarises the same performance measurements numerically by reporting the **AUC** of each method for each set of initial correspondences. All methods perform best using N-ICP [27] for initialisation, with the exception of R3DS Wrap 3, which achieved its best results using None. This slight improvement may be a consequence of the coarse-to-fine strategy being able to converge to a more optimal solution. Ren et al. [155] achieve better results using no correspondence initialisation compared

to Pruning [191] or SEG [100], or just using **WKS** [13] descriptors. In all cases, using SEG [100] produced better results than using Pruning [192].

Table 6.2: (a) The **AUCs** are calculated from the cumulative error curves of the initial correspondence methods reported in Fig. 6.6. (b) The resulting **AUCs** of methods using each set of initial correspondences—reported in Fig. 6.7. Note, Pruning [192] reports an average of 7.4250% vertex correspondences.

(a)

Initialisation method	AUC
N-ICP [27]	0.8736
Pruning [192]	0.6151
SEG [100]	0.5323

(b)

Method	N-ICP [27]	Pruning [192]	SEG [100]	None
R3DS Wrap 3	0.8763	0.5869	0.6050	0.8837
Ezuz and Ben-Chen [67]	0.8829	0.5430	0.5891	N/A
Dyke et al. [55]	0.8771	0.5812	0.6290	N/A
Ren et al. [155]	0.9015	0.5480	0.6347	0.7609

6.4.2 Method comparison

Hereafter, methods initialised with N-ICP [27] correspondences are qualified by an asterisk (*). Comprehensive results—measured by **AUC**—of running each method on each scan—grouped by pose or by internal material—are reported in Table 6.3. Error curves that complement Table 6.3 are given in Figs. 6.8 & 6.9. Scans that are members of a given column were examined collectively. FARM+ achieved the highest overall accuracy of any method using a semi-automatic approach. Of the fully-automatic methods, [155] performed the best. On average all methods performed best on the indented pose. The sporadic non-isometry, limited deformation, and slight topological change may have helped. Conversely, twist was the most challenging pose. The most challenging internal material was the risotto. This is likely to be due to the subtle variations induced on local geometry and how the risotto grains affected the way the rabbit bent in the distinct poses. However, excluding the results of [159] and [200], all methods achieve a particularly high accuracy on the stretch pose. Although there is a large amount of non-rigid motion, the topology remains consistent, and the non-isometric deformation is relatively simple.

Table 6.3: The total area under the curve of scans grouped by the type of pose exhibited, scans grouped by internal material, and the overall performance of each method over all scans is reported. The method that achieved the best results in each configuration is emphasised in **bold**, and the second best in *bold italics*. In the final two rows the mean and standard deviation of each column is reported.

Method	Pose					Internal material				Overall
	Twist	Indent	Inflate	Stretch	Couscous	Risotto	Chickpea			
R3DS Wrap 3	0.9025	0.8129	<i>0.9005</i>	0.9418	<i>0.9265</i>	0.8071	0.9277	0.8837		
Ezuz and Ben-Chen [67]*	0.8730	0.8830	0.8689	0.9195	0.8944	0.8591	0.8979	0.8829		
Rodolà et al. [159]	0.5249	0.6247	0.6136	0.5778	0.5840	0.5570	0.6170	0.5863		
Ren et al. [155]*	0.8708	<i>0.9004</i>	0.9070	<i>0.9420</i>	0.8956	0.8950	<i>0.9122</i>	0.9015		
Bouaziz and Pauly [27]	0.8620	0.8713	0.8646	0.9087	0.8842	0.8579	0.8813	0.8736		
Li et al. [109]*	0.8598	0.8719	0.8651	0.9088	0.8822	0.8581	0.8816	0.8733		
Vestner et al. [200]	0.6369	0.7610	0.6827	0.5521	0.7660	0.5966	0.6695	0.6692		
Dyke et al. [55]*	0.8616	0.8796	0.8712	0.9058	0.8886	0.8604	0.8850	0.8771		
FARM+	0.8999	0.9338	0.8780	0.9438	0.9281	0.9111	0.8992	0.9113		
Mean	0.7636	0.8023	0.7967	0.7932	0.8126	0.7623	0.8008			
Std.	0.1698	0.1318	0.1307	0.2040	0.1484	0.1581	0.1551			

R3DS Wrap 3 performed unexpectedly poorly on the indent pose, with the worst performance of any method on scan05. Rodolà et al. [159] performed poorly in most scenarios, especially in cases where a pose exhibited greater non-isometry. For all poses, except inflate, FARM+ converges to 100% correspondence within the smallest geodesic error. Ren et al. [155] performs the worst on the twist pose, this is understandable as the fused geometry caused by twisting contradicts the assumption of continuity.

As shown in Fig. 6.10, overall, FARM+ performed the best of all methods. Ren et al. [155] performed the best of all the fully-automatic approaches. The performance of [55] is comparable to [109], with only a slight improvement over the initial N-ICP [27] correspondences.

6.5 Conclusions

The problem of establishing correspondence between shapes with different internal materials in challenging poses was considered. A new dataset was created with high quality texture-based ground-truths. The resulting accuracy achieved by using different correspondence initialisation techniques was investigated. In this scenario, N-ICP [27] was discovered to perform the best for correspondence initialisation. Of all the methods, FARM+ achieved the greatest accuracy using a semi-automatic approach, while [155] performed the best of the fully-automatic approaches.












Further work would include an extended investigation into the performance of different shape descriptors and initial correspondence techniques on this dataset.

For many of the methods evaluated, it is unclear how to best optimise the parameters. It was possible to incrementally change the parameters to achieve better results, however this is a time consuming strategy and not necessarily possible in cases where ground-truths are not available. Further investigation into optimal parameters on a range of benchmark datasets is required to give a greater overview of what parameters should be selected in a given scenario.


6.6 Summary

In this chapter, deformations in present datasets were identified to be highly constrained. A dataset that enables the evaluation of complex deformation was proposed. To produce an initial set of correspondences, three methods were selected. The initial correspondence produced by each initialisation method was evaluated without any

subsequent refinement. The influence of each set of initial correspondence on the final results was also evaluated by using them to initialise four refinement methods. Numerous registration and shape correspondence algorithms were benchmarked on the dataset. Finally, the results of each method was reported and compared. The next chapter (Chapter 7) proposes a novel dataset that investigates different aspects of non-rigid correspondence evaluation to those examined in Chapters 5 & 6.

	Twist	Indent	Inflate	Stretch
Couscous	 scan01	 scan02	 scan03	
Risotto	 scan04	 scan05	 scan06	 scan07
Chickpea	 scan08	 scan09	 scan10	 scan11

(a)



(b)

Figure 6.2: Illustrations of the contents of the benchmark dataset. (a) partial (source) meshes, (b) the full/watertight (target) mesh. Note that the stretch pose with couscous filling was omitted in preliminary stages of scanning as its appearance was similar to the same pose with risotto filling. All scans were filled with 36 fl. oz. (or approx. 1065ml) of grains or beans, except for scans where the shape was inflated where they were filled with 56 fl. oz. (or approx. 1656ml).



Figure 6.3: A photo of the model (in the pose of scan00) used for the dataset with markers.

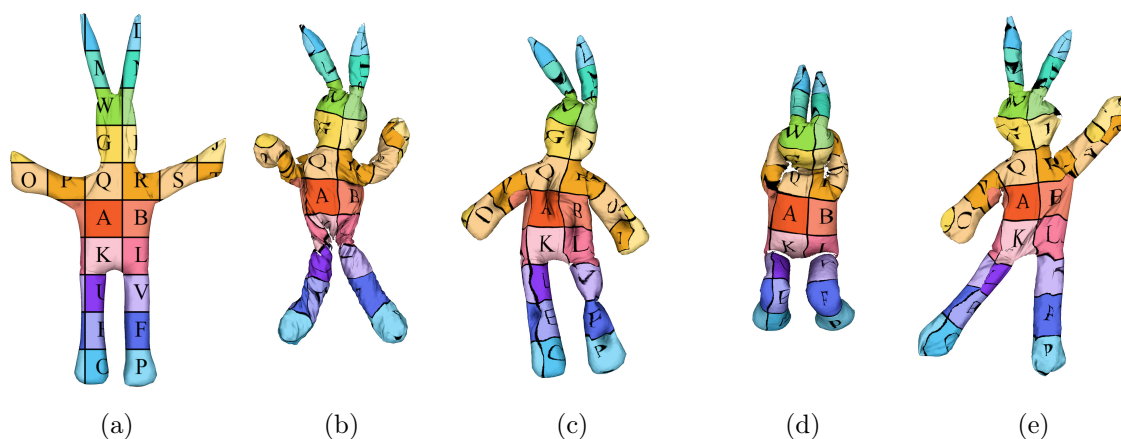


Figure 6.4: Examples of texture transfer using the ground-truth correspondences. (a) Target shape. (b-e) Source shapes. Correspondences were transferred and interpolated using a landmark-based correspondence method [67]. The original texture was projected onto the coronal (frontal) plane of the rabbit in (a).




Figure 6.5: Illustrations of some of the challenges in the dataset. (a) Partial scans (green indicates the boundary). (b) Complex deformations. (c) Missing geometry caused by self-occlusion.




Figure 6.6: Error curves for the methods used to establish an initial correspondence. Note, Pruning [192] only estimated $\approx 7.4250\%$ of correspondences.

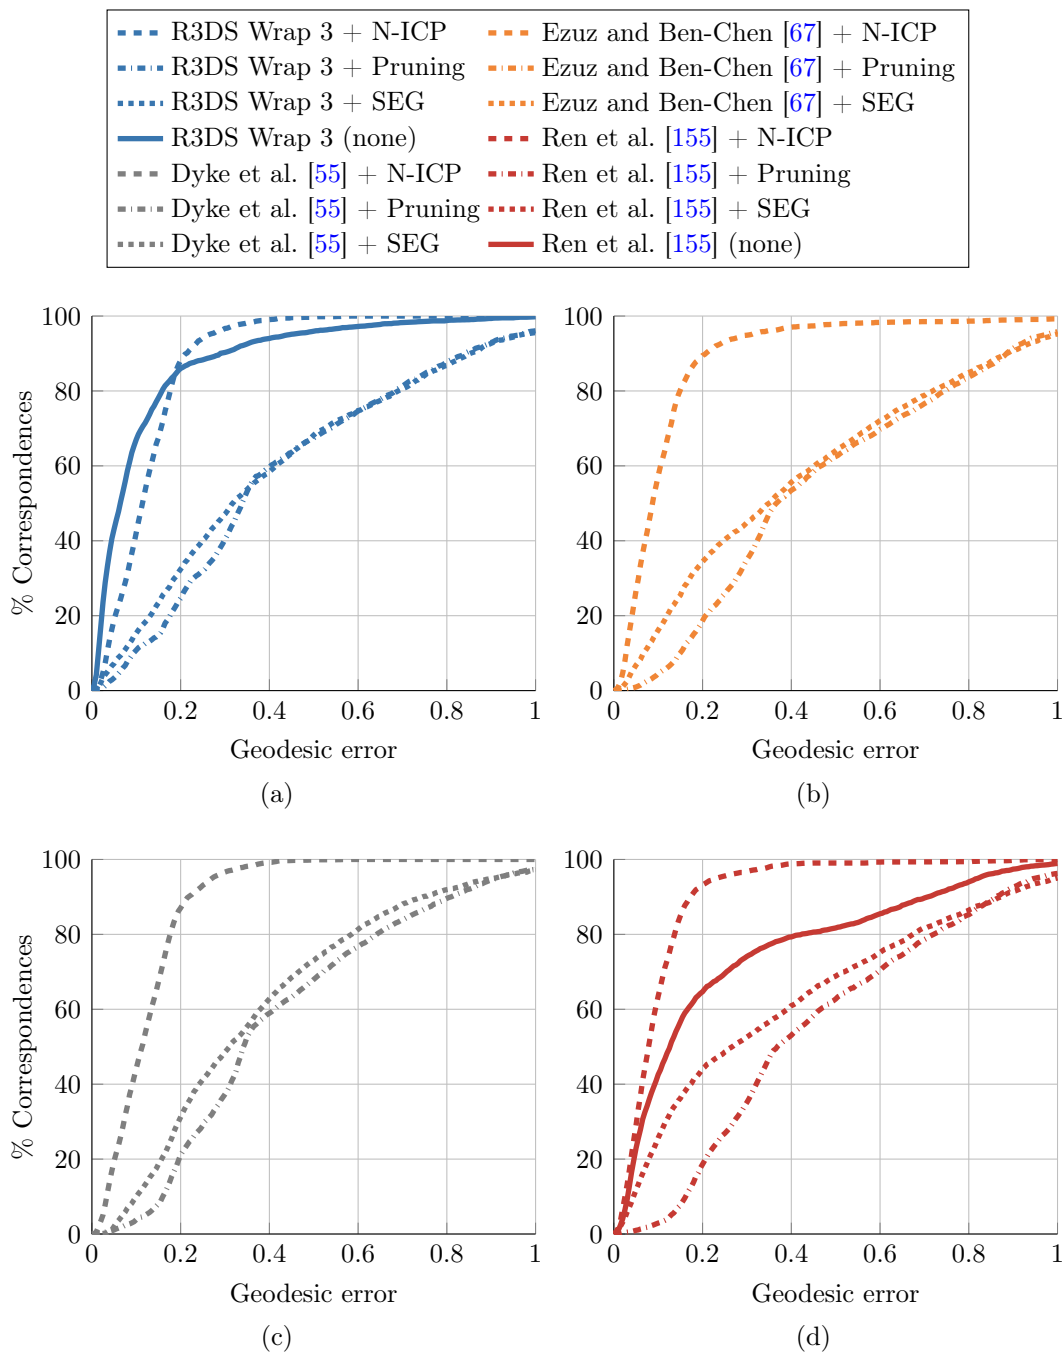


Figure 6.7: Comparison of the error of methods initialised with different initial correspondence techniques (N-ICP [27], Pruning [192], SEG [100], or None). (a) R3DS Wrap 3. (b) Ezuz and Ben-Chen [67]. (c) Dyke et al. [55]. (d) Ren et al. [155].

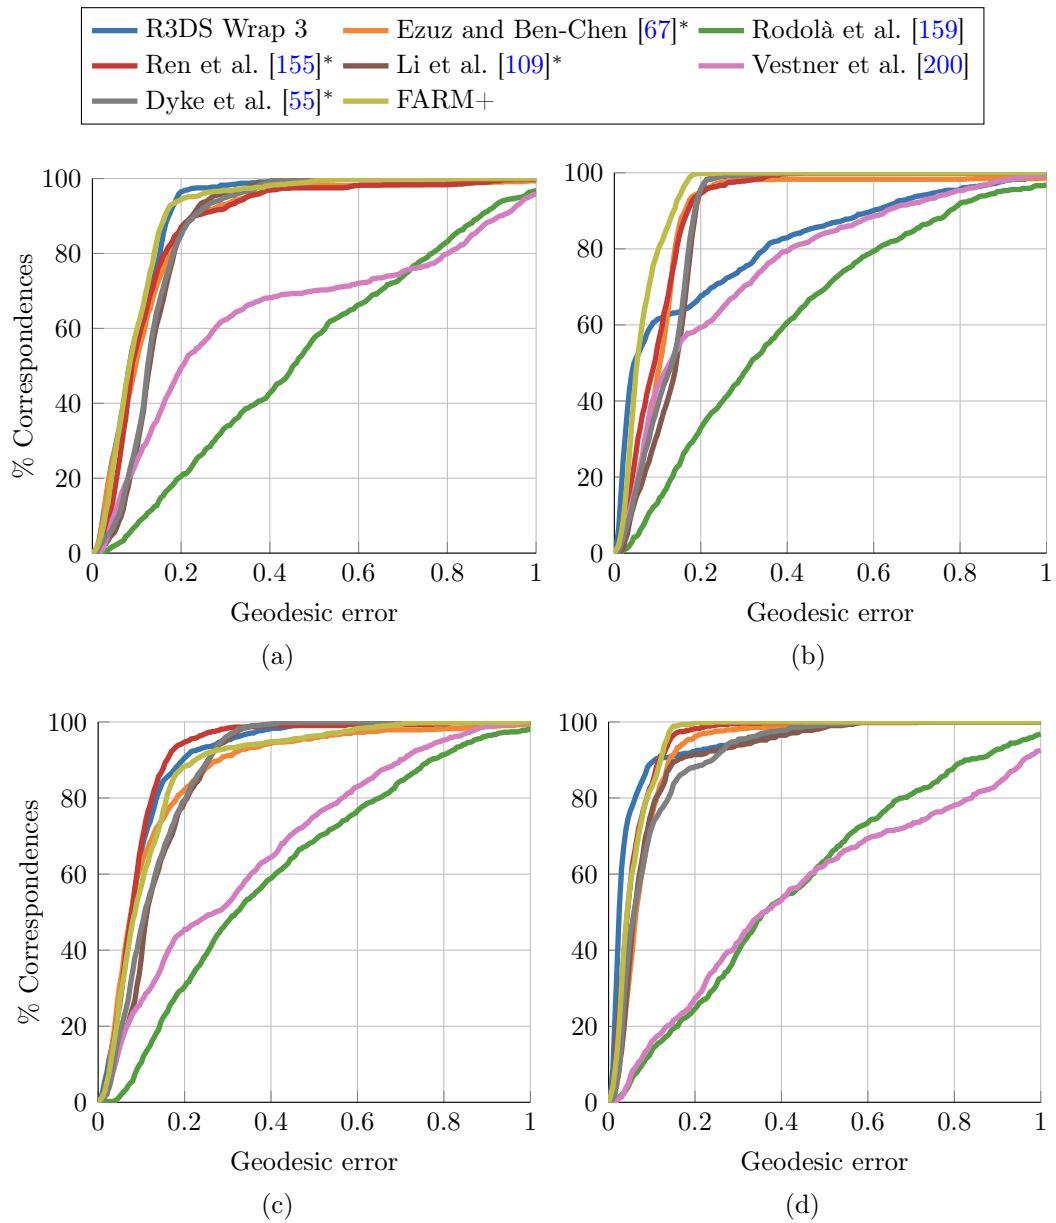


Figure 6.8: Cumulative error curves with scans grouped by the pose exhibited. (a) twist (scan no. 1, 4 & 8). (b) indent (scan no. 2, 5 & 9). (c) inflate (scan no. 3, 6 & 10). (d) stretch (scan no. 7 & 11).

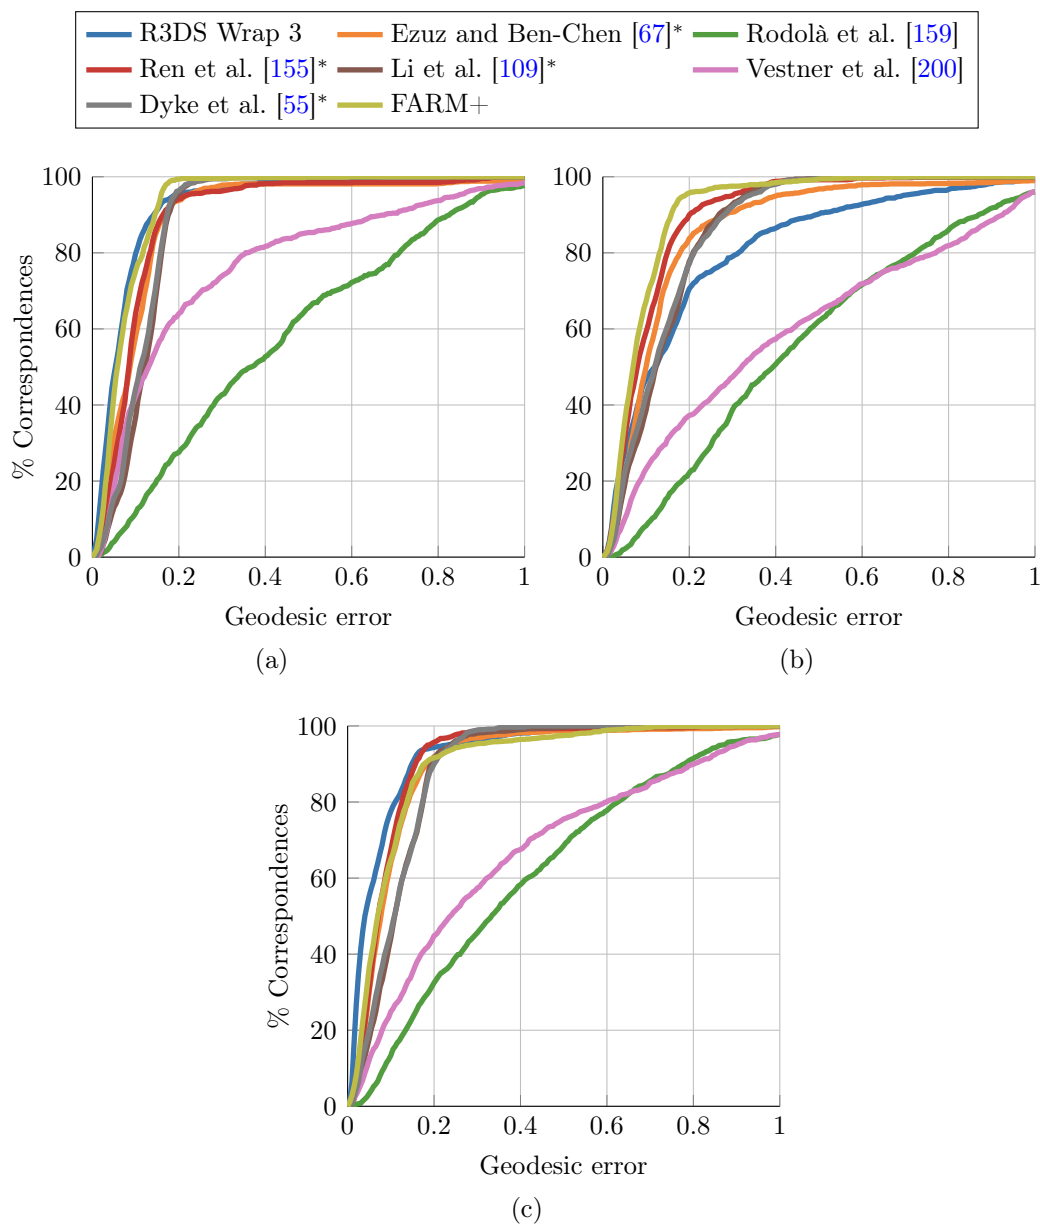


Figure 6.9: Cumulative error curves with scans grouped by the internal material selected. (a) Couscous (scan no. 1, 2 & 3). (b) Risotto (scan no. 4, 5, 6 & 7). (c) Chickpea (scan no. 8, 9, 10 & 11).

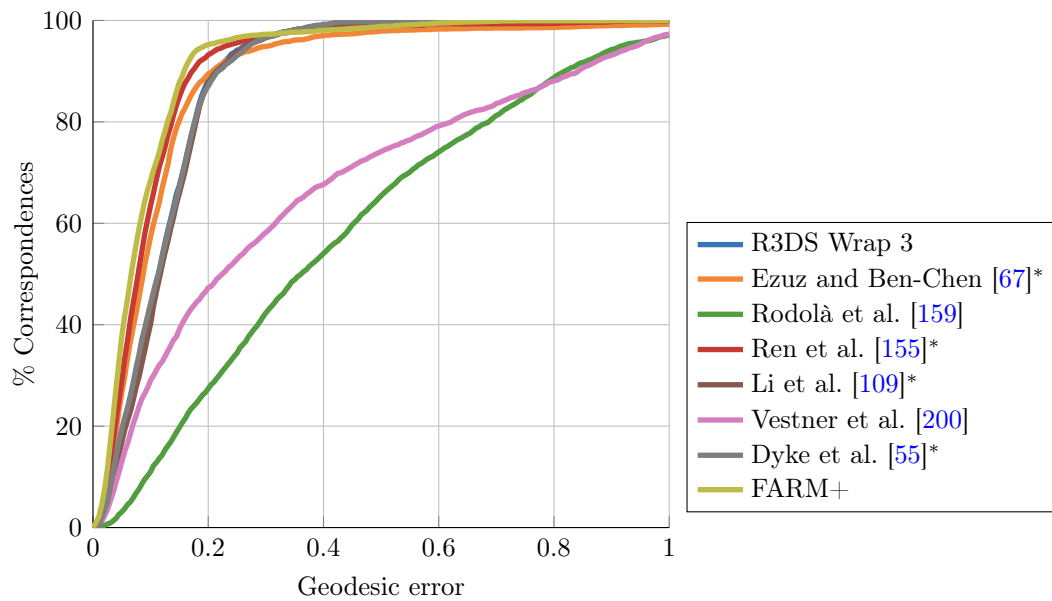


Figure 6.10: The performance of methods over all test sets.

Chapter 7

SHREC’20: Shape correspondence with non-isometric deformations

Overview

In this chapter, a dataset for evaluating non-isometric deformations on quadrupeds is presented. The distinction between the work undertaken in this chapter and Chapters 5 & 6 is highlighted in Section 7.1. Section 7.2 describes the contents of the dataset, as well as specifying the acquisition techniques used to capture each object. Section 7.3 describes the configuration and parameters used for each of the correspondence methods evaluated on this dataset. Section 7.4 describes the measures used to evaluate the performance of methods, including a novel measure of correspondence coverage. In Section 7.5 results are presented and discussed. Finally, Section 7.6 contains concluding thoughts arising from the outputs of this work.

7.1 Introduction

Non-isometric shape correspondence is a problem of great interest. Strictly isometric and near-isometric deformation have been studied extensively, while non-isometric deformation requires further investigation. State-of-the-art methods [74, 55, 88, 66] achieve superior performance in current non-isometric scenarios. There is presently an absence of relevant benchmark datasets for shape correspondence that contain extensive non-isometric deformation.

With respect to the datasets proposed in Chapters 5 & 6, this dataset focuses on a different aspect of non-isometric deformation. In Chapter 5, a range of deformations from piecewise-rigid to non-isometric is considered; however, the degree of

non-isometry is limited. In Chapter 6, a new dataset is created with a focus on complex poses and variations of inner materials that induce non-isometric deformation. In contrast, the poses in the dataset proposed in this chapter are relatively simple. The primary focus of this work is on the wide variation in appearance amongst a range of quadruped animals. The other approaches required costly equipment and do not leverage the use of existing data. Therefore a disproportionate amount of time is required to increase the size of those datasets; whereas, this dataset is compiled using a scalable approach.

Additionally, in real-world scenarios where real objects are scanned, existing capturing techniques induce natural geometric errors (*e.g.*, noise, self-occlusions, and fusion between parts—causing topological changes). The limitations and errors exhibited vary according to the particular scanning technique employed. As in Chapters 5 & 6, most benchmark datasets consist of scans from one scanning device. When evaluating a method’s performance this makes it unclear how well a method’s performance may transfer to other technologies. Recently, Melzi et al. [132] published a dataset that sought to address the issue of incompatibilities between meshes, which arise when working with scans from multiple sources.

Existing datasets that contain a subset of quadruped shapes [187, 182, 32] provide ground-truth correspondences between different shapes of the same class, but not different mammals, limiting the degree of non-isometric deformation that can be quantitatively evaluated. In the case of Bronstein et al. [32], where deforming templates are used, no correspondence information is provided between different templates, so the deformation is near-isometric. The most relevant dataset to this work, by Kim et al. [95], combines three existing datasets [8, 32, 111]. A subset of the shapes are animals, for which a volunteer selected 21 corresponding points on each shape. The dataset contains 51 quadruped shapes. Meshes contain between approximately 3,000 and 56,000 faces. Labels are very sparse, and the expertise of the volunteer is unclear.

For this dataset, a small database of quadruped shapes has been compiled. Establishing correspondence between animal poses is a pertinent challenge that—with shapes exhibiting extreme non-isometries—is not currently considered. As discussed in Section 3.6, most existing datasets address the problem of correspondence between humans, which has quite limited applications. From the perspective of comparative anatomy, being able to also establish correspondence with other mammalian vertebrates—focused on here—may be considered a generalisation of the human correspondence problem.

Although topological variations exist, due to the common ancestry of tetrapod mammals, many parts are considered *homologous structures*, in other words to correspond. Fig. 7.1 illustrates the homologous region of the hind leg between quadrupeds. The hind leg of all tetrapod mammals consist primarily of a femur, tibia, fibula, and metatarsal; therefore, establishing a valid correspondence is possible.

Methods capable of accurately finding correspondences between different mammals enable further avenues of research, *e.g.*, statistical, behavioural analysis [61], and generative models [224]. For zoologists that use morphometrics—the study and development of techniques for the quantitative measurement of organisms—sparse, manually placed correspondences between animals are required to conduct statistical shape analysis.




Figure 7.1: A simple example of a homologous part between quadrupeds. The coloured area represents the matching part for the hind leg. Whilst inter-species, it is possible to intuitively recognise that these areas correspond.

Contribution The main contributions of this work are as follows:

- Generation of a novel dataset of quadrupeds with sparse ground-truth correspondences labelled by three specialists.
- Development of a new measure to evaluate the coverage of correspondences on a shape's surface—discussed further in Section 7.4.
- Systematic evaluation of the performance of a selection of recent shape correspondence methods, with additional quantitative insights into performance from the proposed novel measure.

7.2 Dataset

For this work a set of synthetic models and real-world scans of 3D shapes has been identified. Specifically, the dataset contains only four-legged animals, and a complementing set of ground-truth correspondences has been established. The contents of the dataset is illustrated in Fig. 7.2.




Figure 7.2: Illustrations of the contents of the benchmark dataset simplified to consist of a maximum of 100,000 faces.

Shapes have been amended to remove major errors such as self-intersecting faces and handles which cause erroneous high genera. Ground-truth correspondences were acquired by asking specialists in geometry processing and animal studies to label the shapes manually using a bespoke labelling tool (see Fig. 7.3).

Because the dataset includes real-world scans, many of the shapes contain geometric inconsistency and topological change caused by self-contact. The real-scans also contain natural noise, varying triangulation and self-occluded geometry. Some

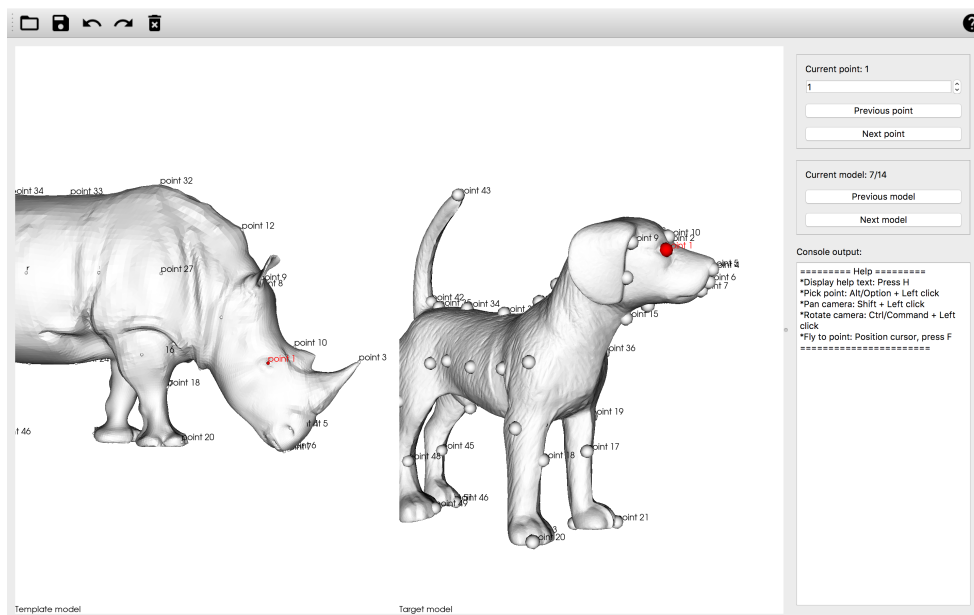


Figure 7.3: A screenshot of the software used by specialists to annotate corresponding points between shapes. The rhino on the left was initially labelled, and used as a reference for subsequent animals.

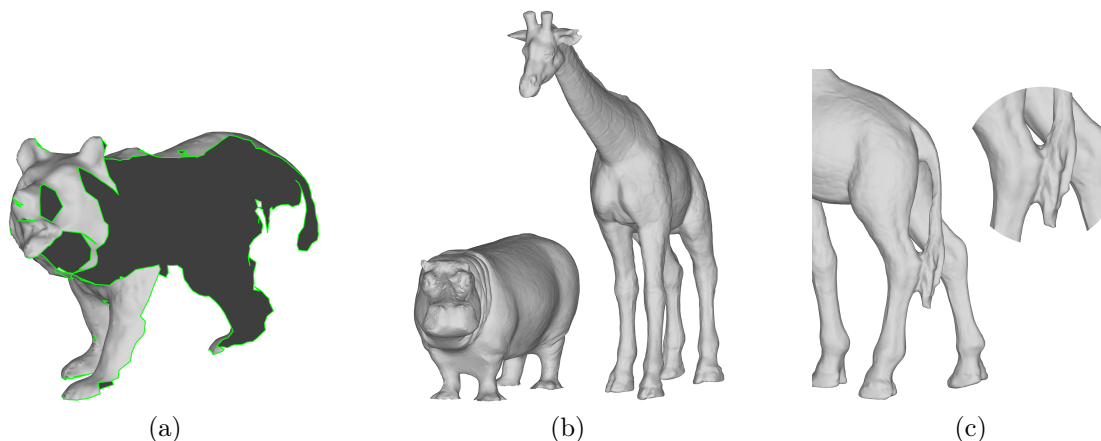


Figure 7.4: Illustrations of some of the challenges in the proposed dataset (a) partial scans (green indicates the boundary), (b) significant non-isometric deformations between pairs of models, and (c) topological inconsistencies: inherent to the original object or caused by scanning limitations.

examples of challenging cases are shown in Fig. 7.4.

The dataset contains 14 models that have been acquired using a variety of techniques (see Table 7.1). Because the dataset is limited to quadruped mammals, many regions share a similar shape or function, it is therefore possible to establish correspondences between homologous loci with a reasonable degree of accuracy. While

the dataset size might initially be considered to be quite small, for the purposes of computing and evaluating corresponding pairs, there are $P(14, 2) = 182$ permutations of shape pairs—or 149 pairs when excluding full-to-partial pairs. In the benchmark experiment, participants were asked to complete a subset of these pairs comprising of matching pairs of full-to-full and partial-to-full models (pairs are listed in Table 7.2).

The ground-truths for this dataset are acquired using the originally sourced meshes. Three specialists labelled corresponding points on each shape based on a template shape that had initially been labelled with markers. For each point, multiple experts propose a correspondence on the surface and a consensus was found by selecting the medoid point. Approximately 50 marker positions were initially selected on the rhinoceros. The rhinoceros was selected as the template since, although it was reconstructed from a multi-view camera array, the shape was subsequently corrected by a professional CGI artist.

For the benchmark, where models have an exceedingly high triangle count, the mesh is simplified to 100,000 triangles. Participants could alternatively submit results using a low-resolution version of the meshes with 20,000 triangles that were also made available.

Ground-truth correspondences were not made available to participants and were solely reserved for evaluative purposes.

Information about the dataset presented here, including how to access it, can be found in the Cardiff University data catalogue, where the dataset has been split into two parts based on the licenses associated with the data (doi: [10.17035/d.2020.0112373427](https://doi.org/10.17035/d.2020.0112373427) (Sketchfab), doi: [10.17035/d.2020.0112716358](https://doi.org/10.17035/d.2020.0112716358) (AIM@SHAPE)).

7.2.1 Test sets

Pairs of scans were carefully selected to ensure the non-isometry present in each test set gradually increased. A description of the contents of each test set may be found in Tab. 7.3.

7.2.2 Initial correspondences

For many shape correspondence and registration algorithms, a sparse set of correspondences is required for initialisation. A set of high quality sparse correspondences enables subsequent automatic refinement of the estimated non-rigid deformation. However, a poor set of initial correspondences may cause the algorithms to fail. For the

purposes of establishing correspondences automatically, it is important to select a robust initial correspondence strategy.

To produce a set of candidate correspondences, **SHOT** signatures [194] at two radii (2% and 5% of the square root of the total triangle area) and *improved wave kernel signature* (IWKS) [112]—a spectral descriptor—were examined, as well as a combination of **SHOT** and IWKS used together. **SHOT** was found to produce the most correct correspondences. A spectral pruning method proposed by Tam et al. [192] was used to remove noisy candidates and produce a set of globally consistent correspondences. For this method to work optimally, input geometry must be locally isometric; however, this was rarely the case in the dataset. Due to limitations of memory and computation time, correspondences were computed with the default parameters, with the exception of $K = 5$ (which specifies how many initial correspondences are found for each point in the source mesh) and $d = 0.25$ (which corresponds to the local neighbourhood size in diffusion pruning).

An illustration of the initial correspondence achieved by each descriptor is given in Fig. 7.5. In areas that are mainly isometric, **SHOT** correctly recovers correspondences. Due to the poor specificity of IWKS, correspondences were often geodesically inconsistent and therefore pruned by diffusion pruning. In areas, such as the head, which has significant non-isometric deformation, it was not possible to establish correspondence using any method.

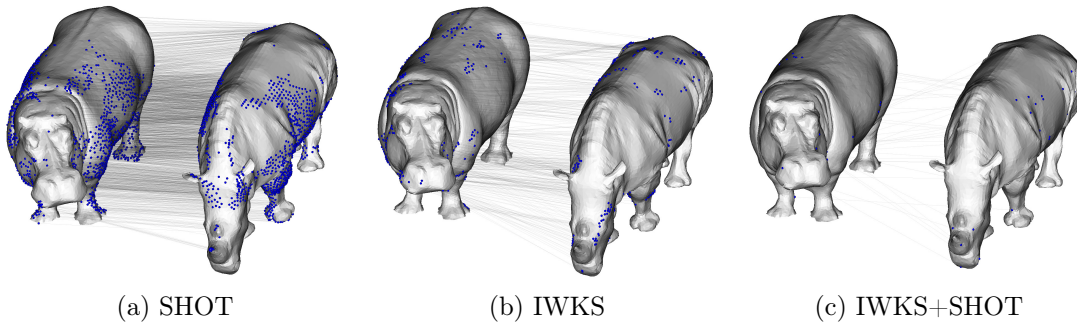


Figure 7.5: An illustration of the initial correspondence computed by matching different descriptors after diffusion pruning.

7.3 Correspondence methods

All the methods examined in this chapter are described in Section 3.5. To avoid duplication, the reader is referred back to the respective method sections for additional

detail, a reference to each method is given in Table 7.4 & 7.5. Specific information about the configurations used in the experiments undertaken here is included.

7.3.1 Baseline N-ICP

For Bouaziz and Pauly [27], the following parameters were used: $w_1=1$, $w_2=1$, $w_3=1$, $w_4=1000$, $iter=100$.

7.3.2 Non-rigid registration under anisotropic deformations

Dyke et al. [55] is initialised using the pre-computed sparse correspondences.

7.3.3 Robust Non-Rigid Registration with Reweighted Position and Transformation Sparsity

Li et al. [109] requires an initial set of sparse correspondences. Pre-computed correspondences were provided to the participants.

7.3.4 Efficient Deformable Shape Correspondence via Kernel Matching

For Vestner et al. [200], all shapes were re-scaled to have a similar area. All parameters remain as described in the original paper, with the following exceptions $\alpha = 10^{-7}$ and $t = \{500, 323, 209, 135, 88, 57, 37, 24, 15, 10\}$.

7.3.5 Deblurring and Denoising of Maps between Shapes

For Ezuz and Ben-Chen [67], due to memory limitations a subset of 200 correspondences were selected using geodesic-based *farthest point sampling*. For experiments the number of basis functions (k_1 & k_2) were set to 120. The method requires an initial set of landmarks. The pre-computed **SHOT** correspondences with diffusion pruning were used for initialisation.

7.3.6 Partial Functional Correspondence

For the partial correspondence method Rodolà et al. [159], shapes were re-scaled to have surface areas between 1.5×10^4 and 2.0×10^4 . All other parameters remain as per their default, with the exception of $n_eigen = 100$.

7.3.7 Continuous and Orientation-preserving Correspondences via Functional Maps

The region-level correspondence method of Kleiman and Ovsjanikov [100] was used to establish an initial correspondence between regions using the default parameters, with the exception of `numComponentsRange = {10, 9, 8, 7}`. An initial functional mapping was computed using Nogneng and Ovsjanikov [142], with $k_1 = 120$ & $k_2 = 120$. The refinement method proposed by Ren et al. [155] was run for 10 iterations to recover point-to-point correspondences.

7.3.8 CMH Connectivity Transfer

We use the *coordinate manifold harmonics* (CMH) framework proposed in Marin et al. [127] and extended on animals in Melzi et al. [134] to establish correspondences by transferring the connectivity. The method computes a functional map following Nogneng and Ovsjanikov [142] using six hand-placed landmarks as probe functions. The match is recovered by finding the nearest neighbour between the target model and the source model deformed using **ARAP** regularisation, and the source connectivity transferred over the model. This method assumes that the target and source shapes share the same pose and does not use the coherent point drift local refinement as proposed in the original paper.

7.3.9 ZoomOut

For ZoomOut by [133], the following parameters were used: 20 as input and 360 as output for the dimension of the functional map, using an incremental step of 10 and 1,000 samples with farthest point sampling for the correspondence step. As with CMH, the connectivity is transferred and the result is refined using **ARAP** optimisation.

7.3.10 R3DS Wrap 3.4

The same implementation of Wrap 3 [167] is used as described in Chapter 6. The default parameters were left unchanged, and the method was provided with initial pruned correspondences.

Submissions As shown in Table 7.6, most methods submitted results for each test set, except Marin et al. [127] & Melzi et al. [133]. This was mainly because these methods were designed to primarily handle cases where objects have the same genus. In this dataset, test-set 2 does not contain any topological changes.

7.4 Evaluation

The quality of correspondences is evaluated using two measures. First, a standard error metric is employed, calculating the geodesic distance between predicted correspondences and the ground-truth locations, as per Kim et al. [95]. Then the quantity and uniformity of correspondences are measured using a novel measurement technique.

7.4.1 Error measure

This chapter uses the error metric described in Section 4.4.1 in Chapter 4 to measure the accuracy of correspondences.

7.4.2 Surface coverage measure

To further assess the performance of methods we develop a *coverage measure*. The measure is derived by first segmenting a shape’s surface into discrete regions and then summing the area of regions that contain a correspondence, this value is then normalised by the shape’s total area. When few regions contain a correspondence the resulting value will be low; this indicates potentially poor overall correspondence between surfaces. By being able to numerically summarise the quality of a set correspondences, it is possible to gain valuable quantitative insights of a method’s performance over a large dataset. Furthermore, as we demonstrate in this section, by varying the number of regions on a shape, it is possible to gain an even greater understanding of how corresponding points are distributed over the target shape. See Alg. 2 for a detailed description of the implementation used.

Region segmentation A set of seed points S on the target surface are selected using a geodesic-based farthest point sampling strategy. This helps ensure a reasonably well distributed sampling. To obtain discrete regions, a Voronoi segmentation Q is subsequently computed using the initial seeds S . Fig. 7.6 illustrates the segmentation of a shape using successively greater numbers of seed points n .

Algorithm 2: Computation of the coverage measure

Input : n – number of seed points/Voronoi cells
 V, F – vertices and triangles of target surface
 A – area of each vertex in surface V, F
 G – geodesic distance map for the surface V, F
 C – binary list with a value of true for vertices with a correspondence

Output: r – ratio of Voronoi cells containing a correspondence

```

seed ← random_integer(1,|V|) // select an initial random seed point
S ← fps_geodesic(G, seed, n) // sample the farthest points on the surface
Q ← find_closest_point(G, S) // return classification of each vertex to the
    closest point in S
for s ← S do
    // iterate through each sampled point on S
    if any((Q == s) ∩ C) then
        // check if there are any correspondences in a given Voronoi cell
        r ← r + ∑i((Qi == s) · Ai) // sum area of the Voronoi cell
    end
end
r ← r / ∑iAi // normalise r by the total surface area

```

Segmentation density By varying n such that $\{n \in \mathbb{N}^+ \mid n \leq |V|\}$ a reasonably smooth and intuitive performance measure is extracted. When $n = 1$, the shape is unsegmented and a single region covers the whole mesh. If the surface has just one correspondence to any point, the output of Alg. 2 would be $r = 1$ (*i.e.*, 100% coverage for $n = 1$). While when $n = |V|$, the barycentric cell of each vertex is a discrete region. To achieve 100% coverage for $n = |V|$, each vertex must have an associated correspondence.

Precise correspondence When handling correspondences with sub-vertex accuracy, each correspondence is associated with the closest discrete point on the shape’s surface. The barycentric cell of each vertex is treated as the vertex neighbourhood, as shown in Fig. 7.7. Points within a given neighbourhood are assigned to that respective vertex.

Segment weighting Since the initial sampling of seed points does not guarantee that the area of each segmented region is uniform, the regions are weighted by their respective area, *i.e.*, where s identifies a unique seed point/region $\sum_i((Q_i == s) \cdot A_i)$. Here, the segmentation classification Q is a list of indices where i refers to a specific vertex and Q_i is the seed point s that is closest to that vertex.

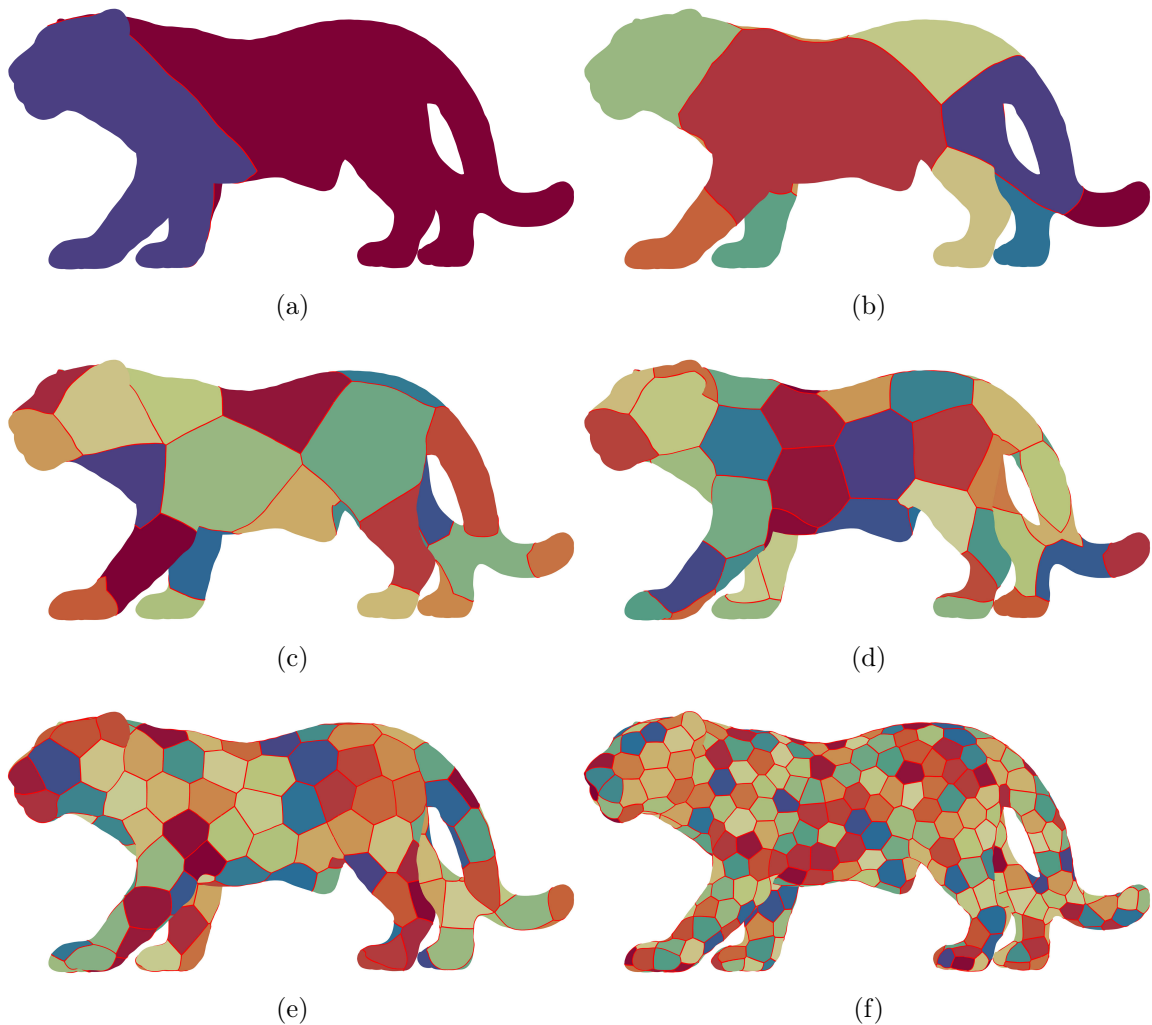


Figure 7.6: Examples of a shape from the dataset segmented using increasingly dense sampling of Voronoi cells. Each sub-figure contains following number of Voronoi cells: (a) 2, (b) 10, (c) 25, (d) 50, (e) 150, & (f) 500.

Computation time Whilst the initial segmentation may be costly to produce, this approach allows the resulting segmentation to be cached and used for further comparison of correspondences with little additional computation. The most costly operation is computing geodesics for the distance map. The complexity of the popular method proposed by Kimmel and Sethian [98] is $O(|V|^2 \log |V|)$.

In Fig. 7.8, an example shape is used to illustrate different ways in which correspondences may be distributed using a set of synthetic points. Fig. 7.9 complements this, demonstrating the response of the coverage measure as the number of segments is varied. It can be seen that the characteristics of the shape of each curve vary by




Figure 7.7: An example illustrating a barycentric cell (the shaded area) for the vertex in the centre of a one-ring neighbourhood. The area of each vertex A_i —used in Alg. 2—also corresponds to the area of the vertex’s barycentric cell.

the type of correspondence computed.

In the case of a bijective mapping, the coverage value will remain at 100% regardless of how fine the segmentation is. Note that this does not mean that the reported correspondences are correct, but that every point on the target surface has a point-to-point correspondence on the source shape. In the case that a part of a shape is not matched (*e.g.*, a leg), the metric will drop off quickly, with no correspondence. Assuming the rest of the shape is successfully matched, the curve should have a gradient equal to zero. For methods that report evenly spaced sparse correspondences, the coverage should remain high until the frequency of the Voronoi cell samples is sufficient to cover areas in-between the sparse correspondences.

7.5 Results

In this section we discuss the performance of correspondence methods with respect to the two measures described in Section 7.4.

7.5.1 Surface coverage

In Fig. 7.10 we measure the coverage achieved by each method that completed all test sets. It is important to note that since a subset of the shapes in the dataset are partial, as the number of Voronoi cells on the surface increases a coverage score of 100% cannot be maintained. The curves of all methods monotonically decrease; this is because all methods report dense correspondences. Therefore, the sparsity of correspondences is not a factor in these results. Based on the curve characteristics discovered in the example of the coverage measure in Section 7.4, Vestner et al. [200]

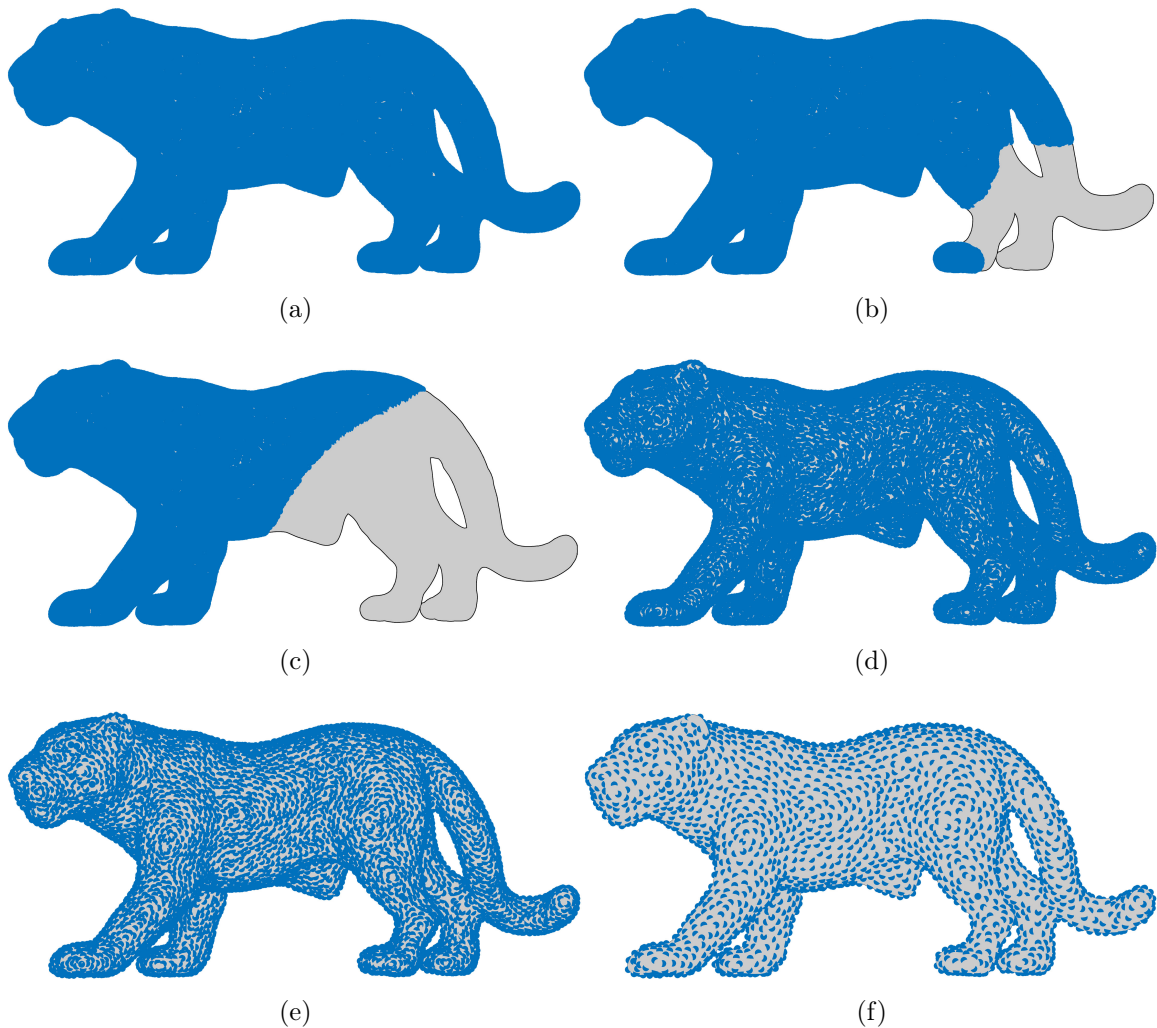


Figure 7.8: Examples of points distributed on the surface of a model from this dataset. (a) Bijective mapping, (b) part mapping (75.0%), (c) part mapping (50.0%), (d) sparse correspondence (50.0%), (e) sparse correspondence (25.0%), and (f) sparse correspondence (10.0%).

exhibit the closest performance to full and dense coverage, with Ren et al. [155] performing second best. This is understandable as both methods promote bijectivity and therefore produce both dense and well distributed correspondences. There is a significant performance gap between the other methods. This may indicate that the other methods do not promote, or do not strongly promote, bijectivity.

Figs. 7.11a & 7.11e report the performance of methods measured by coverage in each respective test set. In test-set 2 (Fig. 7.11c) results for the methods of Melzi et al. [133] and Marin et al. [127] are shown. Both methods perform exceedingly well with a relatively high level of surface coverage. The results for test-sets 1-4 suggest




Figure 7.9: An example of the coverage measure computed using synthetic correspondences on a real mesh.





Figure 7.10: Overall coverage performance of methods that submitted results for all test sets.


that most methods failed to establish correspondences for one or more parts of each target shape.

7.5.2 Geodesic error

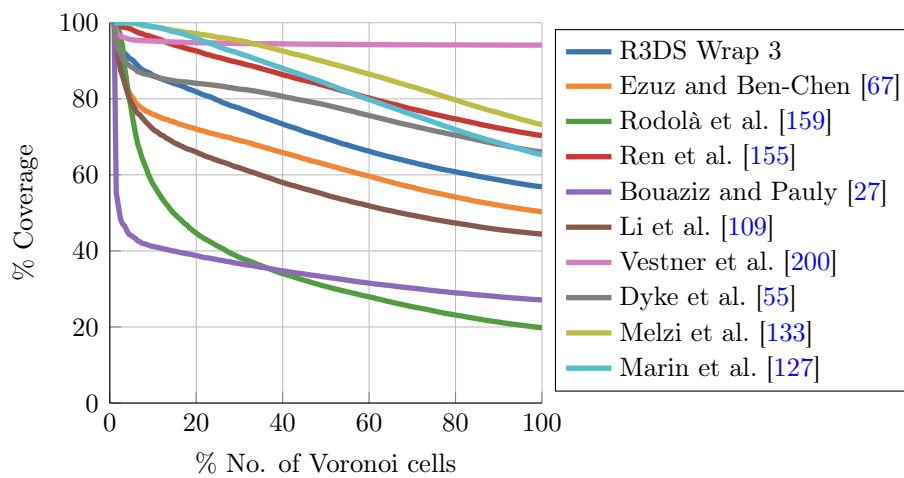
Fig. 7.12 reports cumulative geodesic error curves for methods that have completed all test sets. Table 7.7 complements Fig. 7.12 reporting the AUC of each respective



(a)




(b)




(c)

Figure 7.11: Cumulative error curves for each test set, (a) test-set 0, (b) test-set 1, and (c) test-set 3.



(d)



(e)

Figure 7.11: Cumulative error curves for each test set, (d) test-set 3, and (e) test-set 4.




Figure 7.12: Overall performance of methods that submitted completed all test sets.

method. Overall, most methods appear to perform similarly, with Ren et al. [155] and Bouaziz and Pauly [27] performing the best. However, also taking into account the results of the coverage measure, Ren et al. [155] produce more desirable results with greater coverage of the target shape.

In Fig. 7.13a & 7.13e results for each test set are presented. Results for Melzi et al. [133] and Marin et al. [127] on test-set 2 are shown in Fig. 7.13c. These methods achieve superior correspondence accuracy in comparison to the fully-automatic methods.

Test-set 0 contains only partial-to-full scans. Bouaziz and Pauly [27] perform particularly well, this may be due to shapes having a similar initial orientation, which is important for N-ICP-based methods. Several methods achieve higher levels of accuracy on test-set 2, this may be due to this test-set containing little topological change. With the exception of Ren et al. [155], most methods perform poorly on test-set 4. This is likely to be due to the higher degrees of non-isometry exhibited. The performance of Ren et al. [155] may be due in part to the use of a region-level correspondence method [100] for initialisation, since the correspondence method works particularly well on homogeneous shapes.

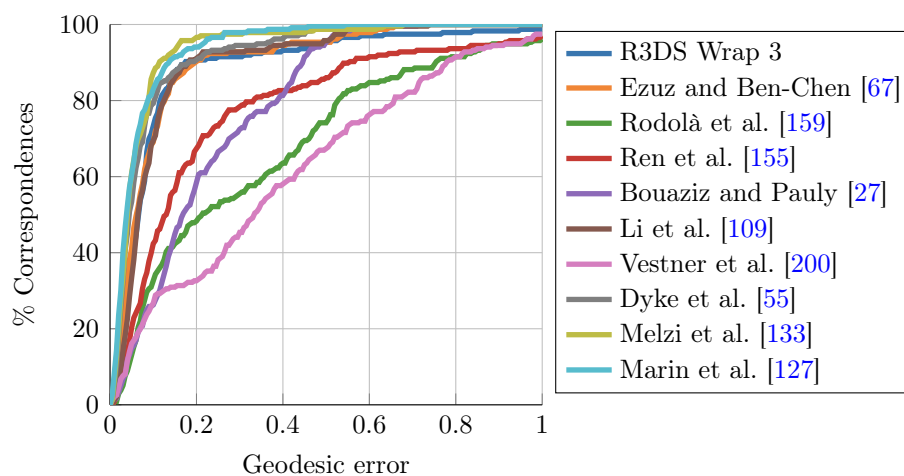






Figure 7.13: Cumulative error curves for each test set, (a) test-set 0, (b) test-set 1, (c) test-set 2.



(d)



(e)

Figure 7.13: Cumulative error curves for each test set, (d) test-set 3 and (e) test-set 4.

7.6 Conclusion

In this chapter, a new benchmark dataset of non-isometric non-rigidly deforming shapes has been proposed to evaluate the performance of shape correspondence methods. To ensure greater accuracy, ground-truth correspondences were established by asking multiple specialists to annotate the animals. The performance of a variety of methods was evaluated using this dataset. Whilst traditional measures of correspondence accuracy are useful, they do not show the full picture of a correspondence method's performance. To address this, a new measure of correspondence coverage has been developed. The coverage measure helps quantitatively indicate the sparsity and distribution of correspondences. Ren et al. [155] was found to achieve the greatest accuracy, as well as a high degree of surface coverage, making it the overall best method in this scenario.

Both Melzi et al. [133] and Marin et al. [127] present semi-automatic methods that achieve superior accuracy compared to the fully-automatic methods. Though not evaluated here, in the paper of Groueix et al. [74]—a data-driven shape correspondence method—the authors make use of an out-of-the-box blend skinning model [224] for training. It would be interesting to see how well this method performs on this dataset where there are animals that have not been observed in the trained model.

As the accessibility of 3D scene capturing tools increases, there is a greater need for high-quality datasets that may be used for benchmarking and training purposes. Restrictive copyrights on existing works make the curation of such datasets challenging. Websites such as Sketchfab and AIM@SHAPE-VISIONAIR Shape Repository provide simple licenses that make it clear what the original creator permits and may enable further datasets to be produced.

The coverage measure formulated here offers important insights into the performance of shape correspondence methods. It can be directly incorporated into many existing benchmarks—such as those mentioned in Section 3.6. In the case of [43], where a dense ground-truth correspondence is known between near-isometrically deforming shapes, it is possible to modify the measure to normalise its output coverage based on an approximation of the total corresponding area. This may help ensure a fair comparison of collections of shape pairs.

This work provides a point of reference on the performance of state-of-the-art methods. The results demonstrate that non-isometric deformation remains a challenging shape matching problem, and further research is still required. This dataset could be extended to include a training facility. The performance of learning-based

methods on unseen animals would be particularly insightful as a measure of generalisation.

7.7 Summary

This chapter identified that many correspondence datasets exist for humans, however the degree of non-isometry is relatively low—making it difficult to assess how suitable a benchmarked method will perform on other datasets. For the purposes of this chapter, it was suggested that the category of quadrupeds is a suitable superset for evaluating greater non-isometries. A collection of quadruped shapes was curated from online sources. Ground-truths were established by experts. Then, numerous state-of-the-art registration and shape correspondence algorithms were benchmarked on the dataset. As well as the standard error measure used in the evaluations in Chapters 4, 5 & 6, a novel measure that quantifies the distribution and sparsity of correspondences was proposed. In the discussions of the results, the new measure was shown to enhance the analysis possible on the correspondence data.

Table 7.1: Summary of models contained in the dataset, as well as a description of the acquisition method used.

Source	Acquisition method	Model name	Original no. of faces	
Sketchfab.com	Metashape by Agisoft ¹	bison*	49,978	
	NextEngine 3D Scanner ²	leopard*	179,680	
	EinScan-S 3D Scanner by SHINING 3D ³	giraffe_a*	2,045,246	
	COLMAP ¹ with 20 Canon G16 cameras	rhino*	64,174	
	Sprout Pro by HP ³	pig*	488,892	
	Metashape by Agisoft ¹	elephant_a*	19,998	
	""	dog*	141,698	
	Recap360 by Autodesk ¹ with an iPhone	hippo#	249,253	
	AIM@SHAPE- VISIONAIR	Synthetic	cow*	14,506
		Unknown	giraffe_b*	18,474
Kemelmacher and Basri [93] ⁴		bear#	31,309	
""		camel_b#	11,698	
Synthetic		camel_a*	19,536	
	Konica-Minolta Vivid 910 ²	elephant_b*	3,024,588	

¹photogrammetry software ²laser triangulation ³structured light ⁴photometric stereo
* full/watertight mesh # partial mesh

Table 7.2: Description of models remeshed to 20,000 (low-resolution) and 100,000 (high-resolution) faces. L_2 - & L_∞ -norm errors with respect to diagonal length of each shapes' bounding box are reported.

Model name	Cleaned no. of faces	High-resolution model			Low-resolution model		
		No. of faces	L_2 error (%)	L_∞ error (%)	No. of faces	L_2 error (%)	L_∞ error (%)
bear	31,309	31,309	0	0	19,999	0.0036	0.0020
bison	49,978	49,978	0	0	20,000	0.0199	0.2298
camel_a	19,528	19,528	0	0	19,528	0	0
camel_b	11,580	11,580	0	0	11,580	0	0
cow	13,866	13,866	0	0	13,866	0	0
dog	14,1698	100,000	0.0027	0.0027	20,000	0.0197	0.2305
elephant_a	19,998	19,998	0	0	19,998	0	0
elephant_b	3,024,548	100,000	0.0059	0.0590	20,000	0.0219	0.1583
giraffe_a	1,325,332	100,000	0.0037	0.0329	20,000	0.0168	0.1184
giraffe_b	18,474	18,474	0	0	18,474	0	0
hippo	153,714	100,000	0.0026	0.0692	20,000	0.0174	0.2702
leopard	179,616	100,000	0.0023	0.0498	20,000	0.0175	0.1443
pig	488,328	100,000	0.0498	0.0214	20,000	0.0147	0.1229
rhino	63,848	63,848	0	0	20,000	0.0147	0.1872

Table 7.3: Summary of shape pairs in each respective test set. All source models in test-set 0 are partial scans and all target models are a full scan. All pairs of models in test-set 1 to 4 are full-to-full scans.

Test-set no.	Source model name	Target model name	Test-set no.	Source model name	Target model name	
0	camel_b	camel_a	2	dog	leopard	
	hippo	rhino		pig	leopard	
	bear	cow		leopard	cow	
	camel_b	cow		dog	pig	
	bear	leopard		cow	dog	
		bear	dog	3	camel_a	cow
		hippo	cow		dog	camel_a
		hippo	dog		rhino	cow
		bear	giraffe_a		pig	elephant_a
		camel_b	elephant_a		bison	elephant_b
1	giraffe_a	giraffe_b	4	rhino	elephant_a	
	elephant_a	elaphant_b		cow	elephant_a	
	camel_a	giraffe_a		elephant_a	giraffe_a	
	giraffe_b	camel_a		cow	giraffe_b	
	cow	bison		dog	giraffe_a	

Reference	Section	Method
[27]	3.5.1	S. Bouaziz and M. Pauly. Dynamic 2D/3D registration for the Kinect. In <i>ACM SIGGRAPH 2013 Courses</i> , SIGGRAPH '13, pages 21:1–21:14. Association for Computing Machinery, July 2013. ISBN 978-1-4503-2339-0. doi: 10.1145/2504435.2504456
[55]	3.5.2	R. M. Dyke, Y.-K. Lai, P. L. Rosin, and G. K. L. Tam. Non-rigid registration under anisotropic deformations. <i>Computer Aided Geometric Design</i> , 71:142–156, May 2019. ISSN 0167-8396. doi: 10.1016/j.cagd.2019.04.014
[109]	3.5.5	K. Li, J. Yang, Y. K. Lai, and D. Guo. Robust non-rigid registration with reweighted position and transformation sparsity. <i>Transactions on Visualization and Computer Graphics</i> , 25(6):2255–2269, June 2018. doi: 10.1109/TVCG.2018.2832136
[200]	3.5.4	M. Vestner, Z. Löhner, A. Boyarski, O. Litany, R. Slossberg, T. Remez, E. Rodolà, A. M. Bronstein, M. M. Bronstein, R. Kimmel, and D. Cremers. Efficient deformable shape correspondence via kernel matching. In <i>International Conference on 3D Vision</i> , 3DV '17, pages 517–526. IEEE, October 2017. doi: 10.1109/3DV.2017.00065
[67]	3.5.9	D. Ezuz and M. Ben-Chen. Deblurring and denoising of maps between shapes. <i>Computer Graphics Forum</i> , 36(5):165–174, August 2017. ISSN 0167-7055. doi: 10.1111/cgf.13254
[159]	3.5.8	E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. <i>Computer Graphics Forum</i> , 36(1):222–236, January 2017. ISSN 0167-7055. doi: 10.1111/cgf.12797
[155]	3.5.10	J. Ren, A. Poulénard, P. Wonka, and M. Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. <i>Transactions on Graphics</i> , 37:248:1–248:16, December 2018. doi: 10.1145/3272127.3275040

Table 7.4: The methods evaluated by the dataset described in this chapter, continued in Table 7.5.

Reference	Section	Method
[127]	3.5.12	R. Marin, S. Melzi, P. Musoni, F. Bardoni, M. Tarini, and U. Castellani. CMH: Coordinates manifold harmonics for functional remeshing. In <i>Eurographics Workshop on 3D Object Retrieval</i> , 3DOR '19, pages 63–70. The Eurographics Association and John Wiley & Sons, Ltd., 2019. doi: 10.2312/3dor.20191063
[133]	3.5.13	S. Melzi, J. Ren, E. Rodolà, A. Sharma, P. Wonka, and M. Ovsjanikov. ZoomOut: Spectral upsampling for efficient shape correspondence. <i>Transactions on Graphics</i> , 38(6):155:1–155:14, November 2019. ISSN 0730-0301. doi: 10.1145/3355089.3356524
[167]	3.5.7	Russian3DScanner. Wrap 3.4, June 2019. URL https://www.russian3dscanner.com/

Table 7.5: The second part of the list of methods evaluated in this chapter.

Table 7.6: A summary of the test sets completed by participants, which are marked with a tick. The absence of a tick indicates the test set was not completed. In the case of Melzi et al. [133] & Marin et al. [127], test-set 2 was the only test set with topologically consistent meshes of the same genus.

Method	Test-set				
	0	1	2	3	4
R3DS Wrap 3	✓	✓	✓	✓	✓
Ezuz and Ben-Chen [67]	✓	✓	✓	✓	✓
Rodolà et al. [159]	✓	✓	✓	✓	✓
Ren et al. [155]	✓	✓	✓	✓	✓
Bouaziz and Pauly [27]	✓	✓	✓	✓	✓
Li et al. [109]	✓	✓	✓	✓	✓
Vestner et al. [200]	✓	✓	✓	✓	✓
Dyke et al. [55]	✓	✓	✓	✓	✓
Melzi et al. [133]			✓		
Marin et al. [127]			✓		

Table 7.7: **AUC** results from the overall geodesic error of each method.

Method	AUC
R3DS Wrap 3	0.708996
Ezuz and Ben-Chen [67]	0.697713
Rodolà et al. [159]	0.606127
Ren et al. [155]	0.747115
Bouaziz and Pauly [27]	0.743213
Li et al. [109]	0.717810
Vestner et al. [200]	0.610986
Dyke et al. [55]	0.726690

Chapter 8

Automatic orbital measures in CT scans using non-rigid registration

Overview

This chapter considers the use of non-rigid registration for the particular application of computing quantitative measures on non-rigid skull models. The problem and contributions of this chapter are clarified in Section 8.1. A description of the proposed registration pipeline is given in Section 8.2. The pipeline is applied to the relatively large dataset described in Section 8.1, the performance of which is examined using two quantitative measures in Section 8.3. Section 8.4 summarises the work undertaken in this chapter and shares some concluding thoughts.

8.1 Introduction

The problem considered here is the measurements of planarity and asymmetry in the orbits of skulls. Skulls typically have two orbits that compose the two eye sockets. This work particularly focuses on accurately establishing correspondence on the frontier between the orbit and the frontal, zygoma, and maxilla bones (see Fig. 8.1). From an anthropometric perspective, this is an interesting problem that will help establish whether asymmetries of the face are caused by genetic or environmental factors. A review of registration methods that seek to align non-rigidly deforming skulls is undertaken in Section 3.9.

There are several key challenges that must be addressed to solve this problem. The main challenge that must be overcome is that of establishing a correspondence between a template and a target skull. Skulls have few geometric features, and exhibit topological changes and non-isometry—combined, these cause many generalised

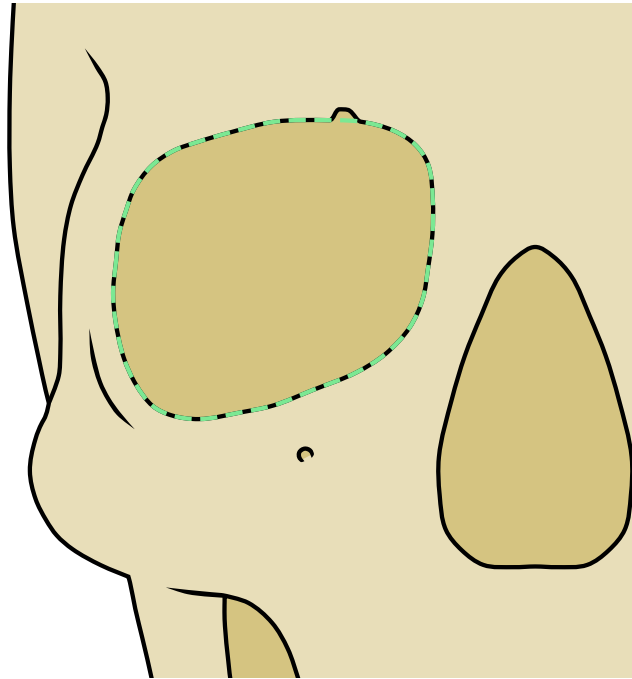


Figure 8.1: An illustrative diagram indicating the approximate location of the orbit. The dashed line represents the orbit on the human skull.

shape correspondence methods to perform poorly. Therefore, a bespoke non-rigid registration method that is suitably robust to such challenges is required. Furthermore, due to the computation time required for registration and shape correspondence algorithms to process high-resolution meshes, low-resolution representations are required. Typically, correspondences are described at a per-vertex level. However, simply using this to transfer correspondences between shapes can lead to a significantly lower accuracy, with template points being projected to their nearest vertex.

Another problem is the representation of orbits. Physically, the orbit is a continuous ring around the eye socket, which can be intuitively imagined but can be difficult to describe in a computationally suitable manner. Parametric representations can be difficult to formulate, with guarantees of all parts of the curve being located on a discrete surface. For this work, a discretised representation of a curve is used to simplify this problem.

To obtain optimal results during surface registration it is important to appropriately condition the registration algorithm used. Relevant local geometric properties (*e.g.*, surface normals and Gaussian curvatures) are considered and used to promote good matches in corresponding regions.

For the layperson, locating orbits can be challenging. Therefore, for the purposes of this work identifying where the orbit actually is, is important. To ensure a rea-

sonable localisation eleven postgraduate orthodontic students were asked to mark 16 points that lie on the frontier of each orbit on ten skulls. A total of 32 points were marked on each skull. One labeller's results were discarded as many of their markers were found, during tests, to be erroneous outliers. This was likely due to insufficient training using the annotation software. Still, the relatively small sample of ten expert labellers is sufficient in this context as experts tend to be consistent.

By fitting cyclic splines to the expert labels at each orbit and then selecting the median spline based on a distance measure (described in Section 8.2.2), ground-truth orbits were produced in a representative manner. Each expert was asked to mark points on the orbit; first the uppermost point, then the lowest point, the point farthest to the right, the point farthest to the left, the midpoint between the uppermost point and right point, and so forth.

Dataset The dataset consists of 1,475 unlabelled scans and ten labelled scans that were annotated by ten experts. The data was captured using *cone-beam computerised tomography* (CBCT). Each scan was cropped to exclude areas surrounding the back of the skull, including the parietal bone, and bones approximately below half way down the neck, or cervical spine. Rather than working directly on the original volume data, a reconstructed representation of the skull surface was used. Each surface was reconstructed using a marching cubes method [122], with each mesh comprising of $\approx 4,000,000$ faces.

For the protection of the individuals in this study, no identifiable craniofacial information that was collected is included in this thesis.

Contributions The main contributions of this work are as follows:

- Development of a set of tools and a framework for transferring information to new skulls, which may be more broadly applied to other collections of meshes.
- An evaluation of the accuracy of the proposed method compared against the performance of a sample of experts.

8.2 Method

The pipeline used to address the problem in Section 8.1 is described here. An overview of the pipeline is illustrated in Fig. 8.2. In the first stage, the mesh of each skull in the database is simplified by discarding disconnected components and reducing the

number of faces, this helps to reduce the time required to process each skull, see Section 8.2.1 for more detail. The second stage (Section 8.2.2) is applied to each of the skulls that has been annotated by experts. For each expert, the positions of their labels are used to produce a closed discretised spline. A distance measure is then used to quantify the similarity between the spline of each expert and another expert’s spline from the same skull. By applying this distance measure to all combinations of expert splines the medoid error can be used to select a spline that is representative of the experts’ labels from the same skull. Due to the computation of the spline, parts of the discretised curve may be distant from the mesh surface; therefore, the discretised curve is projected to sub-vertex locations on the surface of its originating mesh. After this stage, it is possible to transfer the discretised spline from the mesh processed in the previous stage to an unlabelled mesh. To do this, a modified **N-ICP** pipeline is first used to register the meshes. Then, to improve the accuracy of the results, the *mean-value coordinates* (**MVC**) method—a common approach to compute generalised barycentric coordinates—is used to transform the spline to the unlabelled skull. By applying this process to just the labelled skulls, it is possible to establish a suitable template, this is discussed in Section 8.2.3.




Figure 8.2: An overview of the proposed orbit registration pipeline.

8.2.1 Data preparation (preprocessing)

Due to the large size of the dataset that needed to be processed, it was important to save time where possible. To reduce the amount of extraneous information being processed any small disconnected components were removed. For some skulls, the soft tissues in the neck provided sufficient spacing between the neck and skull. This produced discontinuity between the two parts, which was exploited to preserve the skull’s geometry, discarding the neck component. The resulting model was then simplified using a popular mesh simplification technique [70] to reduce the requisite computation time. This process was applied to each model in the dataset.

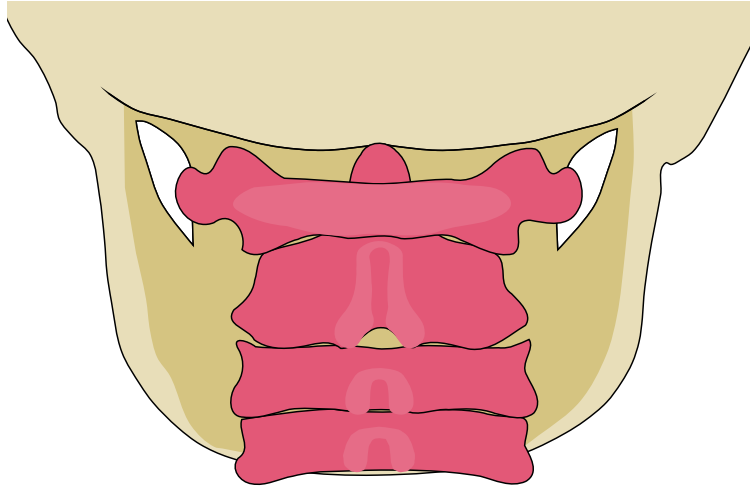


Figure 8.3: An illustrative diagram highlighting the neck bone geometry that is typically excluded due to its separation from the skull, causing a disconnected mesh. The ruby area represents the neck supporting the human skull.

8.2.2 Orbital landmark interpolation

One notable challenge with the landmarks collected from experts was landmark drift along areas of sharp curvature. While participants were instructed to place landmarks at specific points on the orbital frontier (*e.g.*, top, bottom, left, right, etc.), landmarks representing the same location could be placed at a significant distance from each other; however, in many cases the placement was still on, or close to, the orbital boundary, and therefore acceptable for the purposes of this work. To make the landmarks of the experts comparable, markers were interpolated to produce a continuous line that intersected each point smoothly.

Spline generation By fitting a natural cubic spline curve to a sequence of points with the same start and end point, it is possible to produce a closed ring. To ensure landmark points are ordered sequentially, an algorithm designed to solve the Travelling Salesperson Problem is employed. Since there were only 16 points per eye, the computation time is sufficiently quick. Once the spline is computed, a discrete representation containing $\approx 1,000$ points is produced.

Spline distance measure To compare the quality of splines, a one-way measure of error is adopted. First, the nearest neighbours between a source spline \mathbf{s} and a target spline \mathbf{t} are determined using a k -d tree. Denote by \mathcal{N} the set of pairwise correspondences, and $(i, j) \in \mathcal{N}$ means $t_j \in \mathbf{t}$ is the correspondence point of $s_i \in \mathbf{s}$.

Then, for a given point on the source spline, the distance to the corresponding point on the target spline is measured and weighted by the distance to the two adjacent points ($w_{i,i-1}$ and $w_{i,i+1}$) of the source spline point i . The sum of weighted distances is then computed over all source spline points. This may be formally written as

$$E_{\text{spline}} = \frac{1}{2} \sum_{(i,j) \in \mathcal{N}} d_{i,j} \|w_{i,i-1}\| + d_{i,j} \|w_{i,i+1}\|.$$

The definition above is essentially a discretisation of $\int_p d(\mathbf{s}(p)) dp$, where $\mathbf{s}(p)$ is a point on \mathbf{s} and $d(\mathbf{s}(p))$ measures the Euclidean distance from $\mathbf{s}(p)$ to the target. The weights $w_{i,i-1}$ and $w_{i,i+1}$ were required as there is no guarantee that the discretised points on the spline curve are equidistant. This technique is illustrated in Fig. 8.4. The resulting error measured for both the left and right orbit is combined to produce a single error value for the pair of orbits.




Figure 8.4: An illustration of the method used to measure the error between source point i and a discretised spline curve $(\dots, j-1, j, j+1, \dots)$.

Median spline selection For each skull, the error measure described was used to gauge which pair of interpolated orbits best represented the ten experts' selection. The medoid was used due to its innate quality of being highly robust to noise and outliers.

Spline projection Due to the method used to compute the spline curve, there is no guarantee that the curve will lie on the surface of the skull—except where the curve intersects an expert's landmark. For each labelled skull, the pair of median splines are projected onto the surface of the respective skull. However, simply projecting each point to the nearest point on the surface of the skull leads to undesirable discontinuities where a point is projected to the incorrect surface. Due to the




Figure 8.5: An illustration of a correctly projected discretised spline over a one-ring neighbourhood. All adjacent points typically share the same face or are on an adjacent face.

reconstruction method used, this was a common occurrence, which often produced thin walls representing the two sides of the skull bone. To address this problem, an iterative optimisation approach, extending [6], was used.

An efficient point-to-triangle distance technique by Eberly [59] is used to initially project points onto the surface of the mesh. In sequence, each point representing the curve is then checked for continuity by ensuring each adjacent point lies within a given neighbourhood—the neighbourhood size may depend on how dense the discretised spline is, and how coarse the representation of the skull is. Where a point is identified as discontinuous from the previous point, a new segment is created. If subsequent points are continuous with this anomalous point, they are included in the same new segment.

To fix the disharmonious projections, the largest segment is selected—this is assumed to be correct. A point adjacent to the endpoint of the largest segment is projected correctly by limiting the possible projections to the neighbourhood of the current endpoint. Once re-projected, this new point is added to the largest selected segment and is considered to be the new endpoint. This process is repeated until all points form a circular and continuous projected path (illustrated in Fig. 8.5).

Since the spline points tend to be close to the surface already, this process requires few iterations. As described in Chapter 3, the original method by Amberg et al. [6] incorporates an affine matrix-based regularisation term, which is used here. The problem is further constrained by treating the original landmark points provided by the experts as hard correspondences.

8.2.3 Template selection

The aim of this step was to find a skull that would achieve the lowest orbit registration error amongst the experts’ labelled skulls. This skull would be considered the best representation to use as a template for registering the unlabelled dataset. A representative skull was selected from the set of labelled skulls by using the distance between a predicted spline—produced using the registration pipeline—and the ground-truth spline as a measure of error.

Since there were only ten labelled skulls to choose from, each skull was registered to all the other skulls using an extended non-rigid registration pipeline [27]. This gave a total of $P(10, 2) = 90$ pairs of permutations. Once all permutations were registered, the spline of each skull was transferred to all nine targets. The error between the transferred splines and the average splines was computed. The skull that achieved the lowest overall error was then selected as the template (see Fig. 8.6).




Figure 8.6: A registration error matrix illustrating the orbit alignment error of registering all manually labelled skulls against all of the other manually labelled skulls. Skull no. 3 & 5 achieve notably poor results both as the source and as the target scan. In the case of skull no. 3, the spline produced is misaligned. In the case of skull no. 5 the top of the skull is cut off just above the orbit, which is significantly inconsistent with the other skulls.

A more elaborate template can be constructed through the use of techniques such as statistical shape models. However, this was not the main focus of this work, and this approach was found to produce adequate results.

8.2.4 Surface registration

To improve the surface registration results on skulls, with a particular focus on the orbit, extensions to the traditional **N-ICP** implementation [27] are proposed.

Closest point criterion Due to the thin-walled structure of the skulls in the dataset, typical **N-ICP** methods are inadequate, as they often misalign interior and exterior walls. To help distinguish between thin walls, vertex normal information is incorporated. A k -d tree, which for point matching typically contains three-dimensional data for nearest neighbours (*i.e.*, $k = 3$), is used, combining the point's positional information and directional information ($k = 6$). An additional term γ is used to weight the normal information. This additional term is varied from a low to high value (0, 10) throughout the stages of registration. This ensures that normal information is not considered when the source and target surfaces are far apart, but only used to refine the registration results when the surfaces are closer together.

Data term An additional binary term has been added to promote the alignment of points with similar curvature. To ensure this term does not penalise too many matches during the initial steps, only matches that are more than three standard deviations away from the mean are considered to be outliers and removed.

$$E_{\text{point-to-point}} = \sum_{(i,j) \in C} w_i \|\mathbf{x}'_i - \mathbf{y}_j\|_2^2,$$

where,

$$\begin{aligned} \delta_i &= |\kappa_X - \kappa_Y|, \\ \bar{\delta} &= \frac{1}{|C|} \sum_{(i,j) \in C} \delta_i, \\ \sigma^2 &= \sum_{(i,j) \in C} (\delta_i - \bar{\delta})^2, \\ w_i &= \begin{cases} 1 - \frac{\delta_i}{\bar{\delta} + 3\sigma}, & \text{if } \delta_i \leq \bar{\delta} + 3\sigma \\ 0, & \text{otherwise} \end{cases}, \end{aligned}$$

C is the set of correspondences, $(i, j) \in C$ means the source point with existing transformation applied $\mathbf{x}'_i \in X$ corresponds to target point $\mathbf{y}_j \in Y$, and κ_i^X, κ_j^Y are the Gaussian curvatures of points \mathbf{x}_i of surface X and \mathbf{y}_j of surface Y respectively.

Various techniques to incorporate vertex normal information into the energy term were qualitatively assessed to identify an accurate orbit registration.

Orbit transfer While efficient, the proposed registration approach would take an infeasible amount of time to process all 1,475 unlabelled scans at their original density ($\approx 4,000,000$ faces per model). Therefore, to reduce the time required to process the data, each scan was remeshed to a reduced number of triangles. Specifically two sets of registration results were produced using two configurations where meshes comprised of 50,000 faces and 100,000 faces respectively.

Once all the coarse meshes were registered, the orbits were deformed between the two poses using barycentric interpolation. When compared with a nearest neighbour approach, this interpolation method was slightly slower but improved the accuracy of the predicted results when benchmarked using the labelled data as the ground-truth (see Fig. 8.7). Due to the use of coarser representations, this also helped produce smooth orbits. For this work a classical approach was selected [91]. In the experiments undertaken, this method provided superior results to a more complex approach by Lipman et al. [114]. [91] was also significantly faster, as it did not require the additional computation of surface normals for the deformed cage.




Figure 8.7: A comparison of the performance of using a barycentric interpolation method [91] vs. a naïve transportation method—nearest neighbour—on scans remeshed to 50,000 faces.

8.2.5 Quantitative measures

Two measures were incorporated to quantify salient characteristics of each pair of orbits.

Planarity A similar approach to that by King et al. [99] is used to measure how planar a pair of orbits combined. Preliminary results from a *genome-wide association study* (**GWAS**) determined this measure did not produce any significant results.

Asymmetry To more directly determine the difference between the left and right orbits, one orbit is reflected along the axis of symmetry and rigidly aligned to the other. The alignment is optimised using an iterative approach, which incorporates the spline error measure discussed in Section 8.2.2. The formulated optimisation is terminated by a stopping criterion that ensures the two orbits are sufficiently aligned.

8.3 Results

To determine the statistical significance of the results, two-sample *t*-tests were performed to determine whether the data from the ground-truth (*i.e.*, the manually labelled data) and from the estimates belong to the same distribution. Furthermore, paired *t*-tests were used to help find whether the two configurations used to predict the orbit locations belong to the same distribution.

Formally, for comparing the ground-truths and the predicted results, the hypotheses are defined as:

$$H_0 : \mu_{\text{ground-truths}} = \mu_{\text{estimates}}, \text{ and}$$

$$H_1 : \mu_{\text{ground-truths}} \neq \mu_{\text{estimates}}.$$

While for directly comparing the estimates:

$$H_0 : \mu_{\text{estimates}_a} = \mu_{\text{estimates}_b}, \text{ and}$$

$$H_1 : \mu_{\text{estimates}_a} \neq \mu_{\text{estimates}_b}.$$

The estimates of planarity and asymmetry were produced for 1474 of the 1475 scans, the excluded shape was reserved to be the template. The 10 labelled skulls were used to produce the resulting distributions for the ground-truths. Of note is that, when comparing the predicted results to the ground-truths, the observed power is low due to the small sample of ground-truth data. This means that H_0 is likely to be accepted, despite H_1 potentially being correct.

8.3.1 Planarity

The planarity measure, described in Section 8.2.5, is applied to the two configurations used for estimating the location of the orbits as well as the ground-truths. The results of applying the measure are summarised in Fig. 8.8. A planarity value of zero indicates that the orbits lie on the same plane; whereas, anything greater than zero indicates the orbits are non-planar.

Table 8.1: Workings for unpaired parametric two-sample Student t -tests and a paired t -test used to calculate the value of t . This is used to measure the dissimilarity between the orbital planarity results of the ground-truths and estimates. The value t is then compared with the respective critical values to determine the statistical significance of the result.

Group	Sample mean (\bar{x})	$\sum_i^n (x_i - \bar{x})^2$	Variance (s^2)	No. of samples (n)
Ground-truths	0.725450	0.008418	0.0008418	10
Est. (50,000 faces)	0.740499	0.765195	0.0005191	1474
Est. (100,000 faces)	0.737048	0.982426	0.0006665	1474
Combined			s^2	t
G-t/50,000 (unpaired)			0.0005220	-2.075877
G-t/100,000 (unpaired)			0.0006686	-1.413631
50,000/100,000 (paired)			0.0001166	12.2557

G-t/estimates	Value	Estimates/estimates	Value
Significance level (α)	0.05	Significance level (α)	0.05
Df. (ν)	1482	Df. (ν)	1473
Critical value	± 1.961566	Critical value	± 1.961576

Since the data for both the ground-truths and the estimates is known to come from the same population, the variance is treated as the same. Therefore, the number of degrees of freedom is calculated as $\nu = n_1 + n_2 - 2$.

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s^2 \left(\frac{1}{n_1} + \frac{1}{n_2} \right)}}$$

where

$$s^2 = \frac{\sum_i^{n_1} (x_i - \bar{x}_1)^2 + \sum_j^{n_2} (x_j - \bar{x}_2)^2}{\nu}$$

When using triangular meshes with 50,000 faces, we observe that the value t exceeds the lower critical value at $\alpha = 5\%$. Therefore H_0 is rejected and we accept H_1 as true. For meshes with 50,000 faces, we conclude that there is a significant difference between the planarity measured for the ground-truth orbits and the predicted orbits.

However, in the case of meshes with 100,000 faces, the conclusion differs. Since the value t is between the critical values, there is not enough evidence to reject the null hypothesis. We conclude that the difference between the planarity measured for the ground-truth orbits and the predicted orbits is not significant enough.

As the predicted results identically sample the same sources, it is possible to perform a paired t -test. This equates to a one-sample Student t -test calculated on the difference between the measured planarity of the pairs of orbits computed for the same skull. In this case, the number of degrees of freedom is calculated as $\nu = n - 1$ where $n = n_1 = n_2$. Since the value t exceeds the upper critical value, there is sufficient evidence to reject the null hypothesis and to accept the alternative hypothesis. Therefore, we conclude that there is a significant difference between the planarity measured for the predicted orbits with 50,000 faces and the predicted orbits with 100,000 faces.

Since the conclusions for each resolution compared with the ground-truths differ, this may indicate that coarse representations consisting of 50,000 faces are not enough for measuring planarity. This is confirmed by comparing the predicted orbit planarity with 50,000 and 100,000 faces, where the statistical power is high and the value t is found to be beyond the critical value threshold. For 100,000 faces, there is not enough evidence to prove that the distributions differ; however, it is still not possible to declare with confidence that the distributions are equal.




Figure 8.8: A histogram illustrating the planarity results of the predicted orbit registration pipeline using meshes with 50,000 and 100,000 faces. The mean of the ground-truth scans is also shown.

8.3.2 Asymmetry

The asymmetry measure, described in Section 8.2.5, is applied to the two configurations used for estimating the location of the orbits, as well as the ground-truths. The results of applying the measure are summarised in Fig. 8.9. An asymmetry value tending towards zero indicates that the orbits are highly symmetrical, whereas, anything greater than zero indicates an asymmetry.

Table 8.2: Workings for unpaired parametric two-sample Student t -tests and a paired t -test used to calculate the value of t . This is used to measure the dissimilarity between the orbital asymmetry results of the ground-truths and estimates. The value t is then compared with the respective critical values to determine the statistical significance of the result.

Group	Sample mean (\bar{x})	$\sum_i^n (x_i - \bar{x})^2$	Variance (s^2)	No. of samples (n)
Ground-truths	0.052451	0.002576	0.0002576	10
Est. (50,000 faces)	0.122380	0.729652	0.0004950	1474
Est. (100,000 faces)	0.120441	0.774286	0.0005253	1474
Combined			s^2	t
G-t/50,000 (unpaired)			0.0004941	-9.9149389
G-t/100,000 (unpaired)			0.0005242	-9.3589902
50,000/100,000 (paired)			0.0002465	4.7431

G-t/estimates	Value	Estimates/estimates	Value
Significance level (α)	0.05	Significance level (α)	0.05
Df. (ν)	1482	Df. (ν)	1473
Critical value	± 1.961566	Critical value	± 1.961576

Each of the significance tests performed found a notable difference between the asymmetry results of each pair of groups at $\alpha = 5\%$. The value of t for each test was found to be outside of the respective critical values. Therefore, H_0 is rejected and is H_1 accepted. We conclude that there is a significant difference between the asymmetry measured for the ground-truth orbits and the predicted orbits.

From the hypothesis tests performed on the asymmetry results, we can conclude that the results are sufficiently dissimilar. This indicates that the results of the registration for measuring asymmetry in faces are not accurate enough. Since the paired t -test comparing the results of the 50,000 faces to the 100,000 faces also rejected H_0 , this indicates that the representation may be too coarse.




Figure 8.9: A histogram illustrating the asymmetry results of the predicted orbit registration pipeline using meshes with 50,000 and 100,000 faces. The mean of the ground-truth scans is also shown.

8.4 Conclusions

In this work a novel pipeline to automatically estimate the location of a pair of orbits on a skull mesh has been described. The proposed pipeline has a direct application on anthropological research. Two measures that help characterise orbital discrepancies are used to determine the performance of the novel pipeline.

This work identifies that for this particular application 50,000 faces is not enough to produce accurate registration results. Unfortunately, due to the significant amount of time required to process all of the data, it was not possible to run tests with more than 100,000 faces. Ideally, further tests would be conducted to identify an acceptable number of facets for consistent orbit transfer between meshes.

The **MVC** method is incorporated to help compensate for the imprecision caused by coarse representations. There may be scope to improve the accuracy of the obtained results, with further research into the initialisation, template selection, and registration still required. Other issues, such as inconsistent cuts and overfitting, also affected the quality of the results observed. All models had the forehead cut along the horizontal plane manually. In the case of a few models, the cuts were significantly different from the template, which in some cases, caused the forehead geometry of the template to collapse. It may be possible to address this issue by disregarding correspondences that are too far apart, keeping the regularisation of such areas high. Furthermore, the problem of overfitting was observed in the results. Due to the

relaxation of the local regularisation term, in some cases the template skull was over-fitting to geometric errors in the target skull. This registration error would then be propagated into the final estimated spline.

Preliminary tests indicate the use of a novel deep learning technique to predict an initial set of correspondences may help enable faster convergence to the correct solution. Furthermore, for meshes with more than 100,000 faces, a faster **N-ICP** implementation is required to produce results in a feasible amount of time.

No qualitative results are shown in this work due to the sensitive nature of the data. A notable observation about the predicted orbits from this pipeline was that points near the nasal bone were highly varied in both the expert results and predicted results. This was possibly due to the lack of curvature in this area. Also, expert results appeared to have a higher accuracy, yet a lower precision than the proposed pipeline with a greater number of outliers. This could be due to the level of training and the amount of experience the experts had with the annotation software.

8.5 Summary

In this chapter, the automatic measurement of asymmetries localised on the orbit of a skull was identified as a problem that requires attention. A large dataset of human skull meshes obtained using CBCT and reconstructed using a marching cubes approach was described. A bespoke pipeline that is designed to register skull orbits for large datasets was implemented. The large dataset was processed, and the performance of the method was compared with the ground-truths by conducting Student *t*-tests on planarity and asymmetry measures. Additionally, the results obtained using meshes that were simplified to 50,000 faces and 100,000 faces were compared. Based on the results in this chapter, we found that simplifying the skulls to 50,000 faces is insufficient for accurately registering a skull's orbits. Further evaluation is still required to determine whether 100,000 faces is sufficient.

Chapter 9

Conclusions

Overview

The aim of this thesis was to investigate several under-researched issues related to non-rigid registration. In this chapter: The work undertaken in this thesis is summarised in Section 9.1. The key research from each contribution chapter is reviewed in Section 9.2–9.6. Finally, Section 9.7 discusses future work.

9.1 Summary

A systematic literature review into non-rigid registration and shape correspondence techniques was conducted. The investigation highlighted an absence of literature that addresses the problems of large and non-isometric deformation.

The development of a generalised registration pipeline for this little studied problem led to the discovery of the inadequate state of current evaluative benchmarks in the area of non-rigid registration and shape correspondence. Additionally, a review of evaluative benchmarks identified that there was a disproportionate number of anthropomorphic datasets. These datasets were found to be disorganised and, therefore, valuable information that could be extracted by using a structured approach was absent.

This discovery led to the development of three datasets, each designed with distinct goals in mind. While manifesting as theoretical challenges, each of these datasets help measure performance characteristics that must be considered when seeking to address real-world problems. Furthermore, through the development of the first dataset, a need for a descriptive quantitative measure of correspondence sparsity was identified. This was subsequently addressed in Chapter 7 where a measure was proposed and applied to a novel dataset.

In the final contribution chapter, the problem of utilising non-rigid registration for a real medical application was investigated. This led to the development of a bespoke registration pipeline.

9.2 Non-rigid registration under anisotropic deformations

Chapter 4 presented a fully automatic registration method for handling locally non-isometric deformations. Due to its complexity, the problem of non-isometry had received little focus, with most existing techniques assuming deformations to be (near-)isometric. While this assumption is adequate for some scenarios, there are many occasions where deformations are non-isometric.

Based on the evaluations undertaken in Chapter 4, as well as the additional evaluations performed in Chapters 5, 6 & 7, it is important to emphasise the use of the correct method for the correct scenario. The aim of the work in Chapter 4 was to develop a generalised registration pipeline; however, if a problem is well-constrained with suitable training data available, a data-driven approach can potentially achieve better results.

The work aimed to address the non-isometric deformations during registration. A limitation of the proposed implementation is that it relies on **SHOT**, which requires local deformations to be isometric. To improve performance, the use of a descriptor that is compatible with locally non-isometric deformations and has good specificity would be desirable.

9.3 SHREC'19: Shape correspondence with isometric and non-isometric deformations

In Chapter 5, a novel dataset that categorises shape pairs by the type of deformation exhibited was proposed. Real scans were captured of an articulated wooden mannequin and an articulated wooden hand model. To induce different types of deformation, the surface of the mannequin and hand model were covered in textile materials with distinct stretch properties. The type of deformation exhibited was used to form a hierarchical structure for the dataset, which provided insightful information about the performance of a method for different types of deformation. Such a granular organisation had not previously been used in shape correspondence benchmarks.

This dataset was used to compare several state-of-the-art registration and shape correspondence methods. Overall, a registration method [109] that was designed to handle small amounts of locally non-isometric deformation between isometric areas performed the best. However, a data-driven approach [74] was applied to the non-isometric portion of the dataset, and was found to achieve superior results. A baseline method (**N-ICP** [27]) was found to perform well in scenarios where traditional shape descriptors failed due to low-detail homogeneous surfaces.

9.4 SHREC’20: Shape correspondence with physically-based deformations

Presented, in Chapter 6, was a novel dataset that considers a range of challenging poses. Real scans were captured of a stuffed toy rabbit, filled with different internal materials. For each internal material, a scan was captured in each of the poses considered. Most existing shape correspondence datasets only contain human shapes with limited deformation, which are restricted by the internal skeleton. The proposed dataset exploited an object that is capable of extreme poses that are not typically possible in humans, and therefore are typically overlooked.

One unexpected discovery of the work undertaken was the superior performance achieved by using **N-ICP** to compute the initial correspondences rather than other state-of-the-art feature-based approaches. In such scenarios where deformation is limited and surfaces have a good initialisation, traditional **N-ICP** is efficient and can provide adequate results.

The method that performed the best overall was a data-driven approach. This approach was accurate and economic—only requiring a trivial annotation of a template scan to create a deformable model. Ren et al. [155] was found to perform the best of the fully automatic approaches. Two methods performed particularly poorly [159, 200], while both methods have desirable properties for this particular scenario (*i.e.*, Rodolà et al. [159] support partial correspondence, and Vestner et al. [200] support non-isometric deformation) the methods rely on robust feature matching using the **SHOT** descriptor. **SHOT** was found to be sensitive to locally non-isometric deformations; therefore, it was not possible to establish good feature matches.

9.5 SHREC'20: Shape correspondence with non-isometric deformations

Chapter 7 proposed a novel dataset consisting of a range of quadruped mammals with varying physical forms. A range of real and synthetic scans were collected from various online sources. Ground-truth correspondences were established by finding a concordant position for labels placed by three volunteers with expertise in geometry processing and animal studies. Unlike the datasets in Chapters 5 & 6, the shapes selected had simple poses and challenging non-isometries that naturally occur between the corresponding regions of each animal. Furthermore, a novel evaluation metric was proposed to measure the density and distribution of correspondences on a surface.

An evaluation of several state-of-the-art methods was undertaken. Overall, Chapter 7 finds that Ren et al. [155] had the greatest accuracy. The method [155] achieved the second-best coverage score; however, this was still considerably worse than the coverage achieved by Vestner et al. [200]. While [200] achieves a lower accuracy, the estimated correspondence was the closest to bijective. Depending on the application, the inferior accuracy of the method may be an acceptable trade-off for the near-bijective output.

It is possible to further extend the proposed dataset by sourcing new data from existing repositories. When extending a dataset with someone else's work, it is important to consider the restrictions imposed by the original copyright-holder. This can impose restrictions on how the data may be used, and in some cases third-party distribution may not be possible.

9.6 Automatic orbital measures in CT scans using non-rigid registration

In Chapter 8, the problem of extracting information about the orbit of a large database of unlabelled skulls was examined. A representative skull was selected to be a template from amongst a small set of labelled scans. It was then possible to estimate the relative position of orbits on the unlabelled scans using the labelled template through a novel pipeline. Due to the size of the database, it was necessary to reduce computation time through the use of coarse representations of the original scans. To prevent inaccuracies when transferring orbits that had a finer representation, a barycentric interpolation method was used. This was shown to reduce alignment errors of the splines compared to the ground-truths.

A set of hypothesis tests were performed to compare the ground-truths to the estimated measures of planarity and asymmetry for 50,000 and 100,000 faces. Due to the small sample size for the ground-truths, the tests did not have a sufficient power level to confidently conclude that any estimated results belonged to the same distribution as the ground-truths. However, it was possible to say, with a degree of certainty, that the results of planarity measure using 50,000 faces was not enough. In the case of measuring asymmetry, neither 50,000 or 100,000 faces produced sufficiently accurate results. Subsequent paired tests, which had high power levels, confirmed that scans that were remeshed to 50,000 faces were not sufficient for measuring planarity or asymmetry.

The results indicate that further examination of this problem is necessary. The main limitation of the current pipeline is the **N-ICP** implementation that is used. A faster implementation would make it possible to use scans that are less coarse. Additional ground-truth data should be collected to improve the power of the relevant hypothesis tests.

9.7 Future work

The work undertaken in this thesis has helped to progress the area of non-rigid registration; however, these problems are not considered to be solved and more work is necessary. Some potential avenues for future work are suggested in each respective contribution chapter (see Chapters 4, 5, 6, 7 & 8). Further thoughts are given here.

9.7.1 Registration techniques

Further development of techniques to handle non-isometric deformation are required. The evaluations undertaken in Chapters 5, 6 & 7 demonstrate that there is scope to improve upon the current state-of-the-art. Challenging deformations that have not previously been considered were introduced in the dataset proposed in Chapter 6.

For problems in many domains, deep learning has been proven to produce superior results in comparison with hand-crafted methods. While conventional shape descriptors are sensitive to non-isometries, numerous learning-based descriptor methods have been proposed, which show promising results in this area [81]. However, with respect to intrinsic deep learning models, most approaches fail to generalise sufficiently [34].

9.7.2 Datasets

With regards to future work on the topic of this thesis, much work is needed on developing suitable benchmarks that enable the development and evaluation of registration and correspondence techniques for different scenarios. Particularly, further study of how various materials deform and how these might fit or contradict the assumptions of current registration and correspondence methods is required. Also, many shape descriptors have been developed to handle cluttered scenes [89, 194, 209]; however, most existing registration and correspondence methods use data that is reasonably clean. There is little work that evaluates the performance of shape correspondence techniques on cluttered scenes containing non-rigidly deforming shapes [44]. The handling of cluttered scenes is important for many real-world applications, therefore further research into this area is required. With regards to evaluative datasets, there remains scope for a clutter scene database of real scans undergoing varying degrees and types of deformation.

As mentioned in Chapter 3, there is a distinct scarcity of training data for deep learning methods. Such facilities are necessary to ensure parity when comparing the performance of deep learning methods against hand-crafted approaches.

There is scope for further evaluation measures to be developed. An intuitive measurement of accuracy vs. the characteristics of a deformation (*e.g.*, isometry) would provide a greater understanding of how a technique fares with a particular challenge.

9.8 Closing words

This chapter concludes the work that has been conducted for this thesis. Over the previous decade, isometry has been a popular problem in this area. Recently, the problem of non-isometry has received greater attention, and promising progress has already been made. As this work has demonstrated, there are many other challenges for future registration methods to solve. However uncertain the future direction of research in this area may be, research into non-rigid registration over the next decade will hopefully lead to some exciting new developments.

Appendix A

Citation count search terms

The data for this figure was acquired from Thomson Reuters' *Web of Science*.

The following search terms were used: TOPIC: (non-rigid AND shape AND correspondence) OR TOPIC: (non-rigid AND registration) AND TOPIC: (Geometry processing) OR TITLE: (shape AND correspondence)

The results were refined by excluding the following categories: WEB OF SCIENCE CATEGORIES: (TELECOMMUNICATIONS OR NEUROIMAGING OR PHYSICS ATOMIC MOLECULAR CHEMICAL OR AUDIOLOGY SPEECH LANGUAGE PATHOLOGY OR IMAGING SCIENCE PHOTOGRAPHIC TECHNOLOGY OR NEUROSCIENCES OR BEHAVIORAL SCIENCES OR RADIOLOGY NUCLEAR MEDICINE MEDICAL IMAGING OR BIOCHEMICAL RESEARCH METHODS OR ENGINEERING BIOMEDICAL OR BIOTECHNOLOGY APPLIED MICROBIOLOGY OR CARDIAC CARDIOVASCULAR SYSTEMS OR OPTICS OR SURGERY OR DENTISTRY ORAL SURGERY MEDICINE OR ACOUSTICS OR ECOLOGY OR PSYCHOLOGY OR MEDICAL INFORMATICS OR FISHERIES OR PSYCHOLOGY EXPERIMENTAL OR GENETICS HEREDITY OR NANOSCIENCE NANOTECHNOLOGY OR GEOCHEMISTRY GEOPHYSICS OR PHYSICS CONDENSED MATTER OR MARINE FRESHWATER BIOLOGY OR NUCLEAR SCIENCE TECHNOLOGY OR OPHTHALMOLOGY OR BIOPHYSICS OR ORTHOPEDICS OR EVOLUTIONARY BIOLOGY OR PALEONTOLOGY OR FOOD SCIENCE TECHNOLOGY OR LINGUISTICS OR MEDICAL LABORATORY TECHNOLOGY)

Other details: Timespan: All years. Indexes: SCI-EXPANDED, CPCI-S, ESCI.

It should be noted that due to the limited index content and search terms, this does not fully represent all literature in the areas of non-rigid shape correspondence and non-rigid registration. This data is included merely to provide an indication of recent trends in the area.

Bibliography

- [1] J. Achenbach, E. Zell, and M. Botsch. Accurate Face Reconstruction through Anisotropic Fitting and Eye Correction. In *Vision, Modeling & Visualization*, pages 1–8. The Eurographics Association, 2015. ISBN 978-3-905674-95-8. doi:[10.2312/vmv.20151251](https://doi.org/10.2312/vmv.20151251).
- [2] Adobe inc. Mixamo. <https://www.mixamo.com/>, July 2008–2020.
- [3] N. Aigerman and Y. Lipman. Hyperbolic orbifold Tutte embeddings. *Transactions on Graphics*, 35(6), November 2016. ISSN 0730-0301. doi:[10.1145/2980179.2982412](https://doi.org/10.1145/2980179.2982412).
- [4] B. Allen, B. Curless, and Z. Popović. Articulated body deformation from range scan data. *Transactions on Graphics*, 21(3):612–619, July 2002. ISSN 0730-0301. doi:[10.1145/566570.566626](https://doi.org/10.1145/566570.566626).
- [5] P. Alliez. Discrete surfaces. INRIA, April 2017. URL <https://web.archive.org/web/20200809103336/https://team.inria.fr/titane/files/2017/04/discrete-surfaces-new.pdf>.
- [6] B. Amberg, S. Romdhani, and T. Vetter. Optimal step nonrigid ICP algorithms for surface registration. In *Conference on Computer Vision and Pattern Recognition*, CVPR '07, pages 1–8. IEEE, 2007. doi:[10.1109/CVPR.2007.383165](https://doi.org/10.1109/CVPR.2007.383165).
- [7] G. Andrews, S. Endean, R. Dyke, Y. Lai, G. Ffrancon, and G. K. L. Tam. HDFD – A high deformation facial dynamics benchmark for evaluation of non-rigid surface registration and classification. *CoRR*, abs/1807.03354, 2018.
- [8] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. SCAPE: Shape completion and animation of people. *Transactions on Graphics*, 24(3):408–416, July 2005. ISSN 0730-0301. doi:[10.1145/1073204.1073207](https://doi.org/10.1145/1073204.1073207).

- [9] N. Y. Arbel, A. Tal, and L. Zelnik-Manor. Partial correspondence of 3D shapes using properties of the nearest-neighbor field. *Computers & Graphics*, 82:183–192, 2019. ISSN 0097-8493. doi:[10.1016/j.cag.2019.05.011](https://doi.org/10.1016/j.cag.2019.05.011).
- [10] N. Aspert, D. Santa-Cruz, and T. Ebrahimi. MESH: measuring errors between surfaces using the hausdorff distance. In *International Conference on Multimedia and Expo*, volume 1 of *ICME '02*, pages 705–708. IEEE, 2002. doi:[10.1109/ICME.2002.1035879](https://doi.org/10.1109/ICME.2002.1035879).
- [11] M. Attene. A lightweight approach to repairing digitized polygon meshes. *The Visual Computer*, 26(11):1393–1406, November 2010. ISSN 1432-2315. doi:[10.1007/s00371-010-0416-3](https://doi.org/10.1007/s00371-010-0416-3).
- [12] M. Attene. A lightweight approach to repairing digitized polygon meshes. *The Visual Computer*, 26(11):1393–1406, 2010. ISSN 1432-2315. doi:[10.1007/s00371-010-0416-3](https://doi.org/10.1007/s00371-010-0416-3).
- [13] M. Aubry, U. Schlickewei, and D. Cremers. The Wave Kernel Signature: A quantum mechanical approach to shape analysis. In *International Conference on Computer Vision Workshops, ICCVW '11*, pages 1626–1633. IEEE, November 2011. doi:[10.1109/ICCVW.2011.6130444](https://doi.org/10.1109/ICCVW.2011.6130444).
- [14] O. Azencot, É. Corman, M. Ben-Chen, and M. Ovsjanikov. Consistent functional cross field design for mesh quadrangulation. *Transactions on Graphics*, 36(4), July 2017. ISSN 0730-0301. doi:[10.1145/3072959.3073696](https://doi.org/10.1145/3072959.3073696).
- [15] P. Babin, P. Giguère, and F. Pomerleau. Analysis of robust functions for registration algorithms. In *International Conference on Robotics and Automation, ICRA '19*, pages 1451–1457. IEEE, 2019. doi:[10.1109/ICRA.2019.8793791](https://doi.org/10.1109/ICRA.2019.8793791).
- [16] L. Ballan, A. Taneja, J. Gall, L. van Gool, and M. Pollefeys. Motion capture of hands in action using discriminative salient points. In *European Conference on Computer Vision, ECCV '12*, pages 640–653, Berlin, Heidelberg, 2012. Springer. ISBN 978-3-642-33783-3. doi:[10.1007/978-3-642-33783-3_46](https://doi.org/10.1007/978-3-642-33783-3_46).
- [17] G. Barill, N. G. Dickson, R. Schmidt, D. I. W. Levin, and A. Jacobson. Fast winding numbers for soups and clouds. *Transactions on Graphics*, 37(4), July 2018. ISSN 0730-0301. doi:[10.1145/3197517.3201337](https://doi.org/10.1145/3197517.3201337).

- [18] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, September 1975. ISSN 0001-0782. doi:[10.1145/361002.361007](https://doi.org/10.1145/361002.361007).
- [19] P. J. Besl and N. D. McKay. A method for registration of 3D shapes. *Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992. doi:[10.1109/34.121791](https://doi.org/10.1109/34.121791).
- [20] S. Biasotti, A. Cerri, A. M. Bronstein, and M. M. Bronstein. Recent trends, applications, and perspectives in 3D shape similarity assessment. *Computer Graphics Forum*, 35(6):87–119, October 2016. doi:[10.1111/cgf.12734](https://doi.org/10.1111/cgf.12734).
- [21] F. Bogo, J. Romero, M. Loper, and M. J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *Conference on Computer Vision and Pattern Recognition, CVPR '14*, pages 3794–3801. IEEE, June 2014. doi:[10.1109/CVPR.2014.491](https://doi.org/10.1109/CVPR.2014.491).
- [22] F. Bogo, J. Romero, G. Pons-Moll, and M. J. Black. Dynamic FAUST: Registering human bodies in motion. In *Conference on Computer Vision and Pattern Recognition, CVPR '17*, pages 5573–5582. IEEE, 2017.
- [23] N. Bolandzadeh, W. Bischof, C. Flores-Mir, and P. Boulanger. Multimodal registration of three-dimensional maxillofacial cone beam CT and photogrammetry data over time. *Dento maxillo facial radiology*, 42(2):22027087, 2013. ISSN 0250-832X (Print). doi:[10.1259/dmfr/22027087](https://doi.org/10.1259/dmfr/22027087).
- [24] D. Boscaini, J. Masci, E. Rodolà, M. M. Bronstein, and D. Cremers. Anisotropic diffusion descriptors. *Computer Graphics Forum*, 35(2):431–441, 2016. doi:[10.1111/cgf.12844](https://doi.org/10.1111/cgf.12844).
- [25] D. Boscaini, J. Masci, E. Rodolà, and M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *International Conference on Neural Information Processing Systems, NIPS'16*, pages 3197–3205. Curran Associates Inc., 2016. ISBN 9781510838819.
- [26] M. Botsch, L. Kobbelt, M. Pauly, P. Alliez, and B. Levy. *Polygon Mesh Processing*. A K Peters/CRC Press, 1 edition, 2011. ISBN 9781568814261.

- [27] S. Bouaziz and M. Pauly. Dynamic 2D/3D registration for the Kinect. In *ACM SIGGRAPH 2013 Courses*, SIGGRAPH '13, pages 21:1–21:14. Association for Computing Machinery, July 2013. ISBN 978-1-4503-2339-0. doi:[10.1145/2504435.2504456](https://doi.org/10.1145/2504435.2504456).
- [28] S. Bouaziz, A. Tagliasacchi, and M. Pauly. Sparse iterative closest point. In *Proceedings of the Symposium on Geometry Processing*, SGP '13, pages 113–123. The Eurographics Association, 2013. doi:[10.1111/cgf.12178](https://doi.org/10.1111/cgf.12178).
- [29] E. Boyer, A. M. Bronstein, M. M. Bronstein, B. Bustos, T. Darom, R. Horaud, I. Hotz, Y. Keller, J. Keustermans, A. Kovnatsky, R. Litmany, J. Reininghaus, I. Sipiran, D. Smeets, P. Suetens, D. Vandermeulen, A. Zaharescu, and V. Zobel. SHREC '11: Robust feature detection and description benchmark. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '11, pages 71–78. The Eurographics Association, 2011. doi:[10.2312/3DOR/3DOR11/071-078](https://doi.org/10.2312/3DOR/3DOR11/071-078).
- [30] A. M. Bronstein, M. M. Bronstein, U. Castellani, A. Dubrovina, L. J. Guibas, R. P. Horaud, R. Kimmel, D. Knossow, E. v. Lavante, D. Mateus, M. Ovsjanikov, and A. Sharma. SHREC '10: Correspondence finding. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '10, pages 87–91. The Eurographics Association, October 2010. ISBN 978-3-905674-22-4. doi:[10.2312/3DOR/3DOR10/087-091](https://doi.org/10.2312/3DOR/3DOR10/087-091).
- [31] A. M. Bronstein, M. M. Bronstein, U. Castellani, B. Falcidieno, A. Fusiello, A. Godil, L. J. Guibas, I. Kokkinos, Z. Lian, M. Ovsjanikov, G. Patané, M. Spagnuolo, and R. Toldo. SHREC '10: Robust shape retrieval. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '10, pages 71–78. The Eurographics Association, 2010. doi:[10.2312/3DOR/3DOR10/071-078](https://doi.org/10.2312/3DOR/3DOR10/071-078).
- [32] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer-Verlag, 1 edition, 2008. ISBN 0387733000, 9780387733005.
- [33] A. M. Bronstein, M. M. Bronstein, R. Kimmel, M. Mahmoudi, and G. Sapiro. A Gromov-Hausdorff Framework with Diffusion Geometry for Topologically-Robust Non-rigid Shape Matching. *International Journal of Computer Vision*, 89(2):266–286, 2010. ISSN 1573-1405. doi:[10.1007/s11263-009-0301-6](https://doi.org/10.1007/s11263-009-0301-6).

- [34] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *Signal Processing Magazine*, 34(4):18–42, July 2017. doi:[10.1109/MSP.2017.2693418](https://doi.org/10.1109/MSP.2017.2693418).
- [35] B. J. Brown and S. Rusinkiewicz. Non-rigid range-scan alignment using thin-plate splines. In *International Symposium on 3D Data Processing, Visualization and Transmission*, 3DPVT '04, pages 759–765. IEEE, 2004. doi:[10.1109/TDPVT.2004.1335392](https://doi.org/10.1109/TDPVT.2004.1335392).
- [36] M. Chen, C. Wang, and H. Qin. Jointly learning shape descriptors and their correspondence via deep triplet CNNs. *Computer Aided Geometric Design*, 62:192–205, 2018. ISSN 0167-8396. doi:[10.1016/j.cagd.2018.03.002](https://doi.org/10.1016/j.cagd.2018.03.002).
- [37] Q. Chen and V. Koltun. Robust nonrigid registration by convex optimization. In *International Conference on Computer Vision*, ICCV '15, pages 2039–2047. IEEE, 2015. doi:[10.1109/ICCV.2015.236](https://doi.org/10.1109/ICCV.2015.236).
- [38] S.-Y. Chen, L. Gao, Y.-K. Lai, and S. Xia. Rigidity controllable as-rigid-as-possible shape deformation. *Graphical Models*, 91:13–21, 2017. ISSN 1524-0703. doi:[10.1016/j.gmod.2017.02.005](https://doi.org/10.1016/j.gmod.2017.02.005).
- [39] Y. Chen and G. Medioni. Object modelling by registration of multiple range images. *Image Vision Computing*, 10(3):145–155, April 1992. ISSN 0262-8856. doi:[10.1016/0262-8856\(92\)90066-C](https://doi.org/10.1016/0262-8856(92)90066-C).
- [40] P. Cignoni, C. Rocchini, and R. Scopigno. Metro: Measuring error on simplified surfaces. *Computer Graphics Forum*, 17(2):167–174, 1998. doi:[10.1111/1467-8659.00236](https://doi.org/10.1111/1467-8659.00236).
- [41] A. Cohen and M. Ben-Chen. Robust shape collection matching and correspondence from shape differences. *Computer Graphics Forum*, 39(2):555–568, 2020. doi:[10.1111/cgf.13952](https://doi.org/10.1111/cgf.13952).
- [42] É. Corman, M. Ovsjanikov, and A. Chambolle. Supervised descriptor learning for non-rigid shape matching. In *European Conference on Computer Vision Workshops*, ECCV '14, pages 283–298, Cham, 2015. Springer. ISBN 978-3-319-16220-1. doi:[10.1007/978-3-319-16220-1_20](https://doi.org/10.1007/978-3-319-16220-1_20).
- [43] L. Cosmo, E. Rodolà, M. M. Bronstein, A. Torsello, D. Cremers, and Y. Sahillioglu. Partial matching of deformable shapes. In *Eurographics Workshop*

- on 3D Object Retrieval*, 3DOR '16, Goslar, Germany, May 2016. The Eurographics Association. ISBN 978-3-03868-004-8. doi:[10.2312/3dor.20161089](https://doi.org/10.2312/3dor.20161089).
- [44] L. Cosmo, E. Rodolà, J. Masci, A. Torsello, and M. M. Bronstein. Matching deformable objects in clutter. In *International Conference on 3D Vision*, 3DV '16, pages 1–10. IEEE, 2016. doi:[10.1109/3DV.2016.10](https://doi.org/10.1109/3DV.2016.10).
- [45] L. Cosmo, E. Rodolà, A. Albarelli, F. Mémoli, and D. Cremers. Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum*, 36(1):209–221, 2017. doi:[10.1111/cgf.12796](https://doi.org/10.1111/cgf.12796).
- [46] W. R. Crum, T. Hartkens, and D. L. G. Hill. Non-rigid image registration: theory and practice. *The British journal of radiology*, 77 Spec No 2:S140–53, 2004. ISSN 0007-1285 (Print). doi:[10.1259/bjr/25329214](https://doi.org/10.1259/bjr/25329214).
- [47] H. Dai, W. A. P. Smith, N. Pears, and C. Duncan. Symmetry-factored statistical modelling of craniofacial shape. In *International Conference on Computer Vision Workshops*, ICCVW '17, pages 786–794. IEEE, 2017. doi:[10.1109/ICCVW.2017.98](https://doi.org/10.1109/ICCVW.2017.98).
- [48] H. Dai, N. Pears, and W. Smith. Non-rigid 3D shape registration using an adaptive template. In *European Conference on Computer Vision Workshops*, ECCV '18, pages 48–63, Cham, 2019. Springer. ISBN 978-3-030-11018-5. doi:[10.1007/978-3-030-11018-5_5](https://doi.org/10.1007/978-3-030-11018-5_5).
- [49] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor. A minimum description length approach to statistical shape modeling. *Transactions on Medical Imaging*, 21(5):525–537, August 2002. doi:[10.1109/TMI.2002.1009388](https://doi.org/10.1109/TMI.2002.1009388).
- [50] R. H. Davies, C. J. Twining, T. F. Cootes, J. C. Waterton, and C. J. Taylor. 3D statistical shape models using direct optimisation of description length. In *European Conference on Computer Vision*, ECCV '02, pages 3–20. Springer, 2002. ISBN 978-3-540-47977-2. doi:[10.1007/3-540-47977-5_1](https://doi.org/10.1007/3-540-47977-5_1).
- [51] Q. Deng, M. Zhou, W. Shui, Z. Wu, Y. Ji, and R. Bai. A novel skull registration based on global and local deformations for craniofacial reconstruction. *Forensic Science International*, 208(1):95–102, 2011. ISSN 0379-0738. doi:[10.1016/j.forsciint.2010.11.011](https://doi.org/10.1016/j.forsciint.2010.11.011).

- [52] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959. ISSN 0945-3245. doi:[10.1007/BF01386390](https://doi.org/10.1007/BF01386390).
- [53] N. Donati, A. Sharma, and M. Ovsjanikov. Deep geometric functional maps: Robust feature learning for shape correspondence. In *Conference on Computer Vision and Pattern Recognition, CVPR '20*, pages 8589–8598. IEEE/CVF, 2020. doi:[10.1109/CVPR42600.2020.00862](https://doi.org/10.1109/CVPR42600.2020.00862).
- [54] G.-G. Du, C.-L. Yin, M.-Q. Zhou, Z.-K. Wu, Y.-C. Fan, F.-Q. Duan, and P.-B. Zhou. Isometric 3D shape partial matching using GD-DNA. *Journal of Computer Science and Technology*, 33(6):1178–1191, 2018. ISSN 1860-4749. doi:[10.1007/s11390-018-1880-x](https://doi.org/10.1007/s11390-018-1880-x).
- [55] R. M. Dyke, Y.-K. Lai, P. L. Rosin, and G. K. L. Tam. Non-rigid registration under anisotropic deformations. *Computer Aided Geometric Design*, 71:142–156, May 2019. ISSN 0167-8396. doi:[10.1016/j.cagd.2019.04.014](https://doi.org/10.1016/j.cagd.2019.04.014).
- [56] R. M. Dyke, C. Stride, Y.-K. Lai, P. L. Rosin, M. Aubry, A. Boyarski, A. M. Bronstein, M. M. Bronstein, D. Cremers, M. Fisher, T. Groueix, D. Guo, V. G. Kim, R. Kimmel, Z. Löhner, K. Li, O. Litany, T. Remez, E. Rodolà, B. C. Russell, Y. Sahillioğlu, R. Slossberg, G. K. L. Tam, M. Vestner, Z. Wu, and J. Yang. Shape correspondence with isometric and non-isometric deformations. In *Eurographics Workshop on 3D Object Retrieval, 3DOR '19*, pages 111–119. The Eurographics Association, May 2019. ISBN 978-3-03868-077-2. doi:[10.2312/3dor.20191069](https://doi.org/10.2312/3dor.20191069).
- [57] R. M. Dyke, Z. Feng, Y.-K. Lai, P. L. Rosin, D. Guo, K. Li, R. Marin, and J. Yang. Shape correspondence of physically-based deformations. In *Eurographics Workshop on 3D Object Retrieval, 3DOR '20*. The Eurographics Association, 2020. doi:[10.2312/3dor.20201161](https://doi.org/10.2312/3dor.20201161).
- [58] R. M. Dyke, Y.-K. Lai, P. L. Rosin, S. Zappalà, S. L. Dykes, D. Guo, K. Li, R. Marin, S. Melzi, and J. Yang. Shape correspondence with non-isometric deformations. *Computers & Graphics*, 2020. ISSN 0097-8493. doi:[10.1016/j.cag.2020.08.008](https://doi.org/10.1016/j.cag.2020.08.008).
- [59] D. Eberly. Distance between point and triangle in 3D. Technical Report 1, Geometric Tools, Redmond WA 98052, September 1999. URL <https://www.geometrictools.com/Documentation/DistancePoint3Triangle3.pdf>.

- [60] M. Edelstein, D. Ezuz, and M. Ben-Chen. ENIGMA: Evolutionary Non-Isometric Geometry MAtching. *Transactions on Graphics*, 39(4), July 2020. doi:[10.1145/3386569.3392447](https://doi.org/10.1145/3386569.3392447).
- [61] K. Ehsani, H. Bagherinezhad, J. Redmon, R. Mottaghi, and A. Farhadi. Who let the dogs out? modeling dog behavior from visual data. In *Conference on Computer Vision and Pattern Recognition*, CVPR '18, pages 4051–4060. IEEE, 2018. doi:[10.1109/CVPR.2018.00426](https://doi.org/10.1109/CVPR.2018.00426).
- [62] M. Eisenberger, Z. Löhner, and D. Cremers. Divergence-free shape correspondence by deformation. *Computer Graphics Forum*, 38(5):1–12, 2019. doi:[10.1111/cgf.13785](https://doi.org/10.1111/cgf.13785).
- [63] M. Eisenberger, Z. Löhner, and D. Cremers. Smooth Shells: Multi-scale shape registration with functional maps. In *Conference on Computer Vision and Pattern Recognition*, CVPR '20, pages 12262–12271. IEEE/CVF, June 2020. doi:[10.1109/CVPR42600.2020.01228](https://doi.org/10.1109/CVPR42600.2020.01228).
- [64] L. Euler. Recherches sur la courbure des surfaces. *Mem. de l'Academie des Sciences de Berlin*, 16(119-143):9, 1760.
- [65] D. Eynard, E. Rodolà, K. Glashoff, and M. M. Bronstein. Coupled functional maps. In *International Conference on 3D Vision*, 3DV '16, pages 399–407. IEEE, 2016. doi:[10.1109/3DV.2016.49](https://doi.org/10.1109/3DV.2016.49).
- [66] D. Ezuz, B. Heeren, O. Azencot, M. Rumpf, and M. Ben-Chen. Elastic correspondence between triangle meshes. *Computer Graphics Forum*, 38(2):121–134, June 2019. doi:[10.1111/cgf.13624](https://doi.org/10.1111/cgf.13624).
- [67] D. Ezuz and M. Ben-Chen. Deblurring and denoising of maps between shapes. *Computer Graphics Forum*, 36(5):165–174, August 2017. ISSN 0167-7055. doi:[10.1111/cgf.13254](https://doi.org/10.1111/cgf.13254).
- [68] D. Ezuz, J. Solomon, and M. Ben-Chen. Reversible harmonic maps between discrete surfaces. *Transactions on Graphics*, 38(2), March 2019. ISSN 0730-0301. doi:[10.1145/3202660](https://doi.org/10.1145/3202660).
- [69] A. Frome, D. Huber, R. Kolluri, T. Bülow, and J. Malik. *Recognizing Objects in Range Data Using Regional Point Descriptors*, pages 224–237. ECCV '04. Springer, Berlin, Heidelberg, 2004. doi:[10.1007/978-3-540-24672-5_18](https://doi.org/10.1007/978-3-540-24672-5_18).

- [70] M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. In *Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, pages 209–216. ACM Press/Addison-Wesley Publishing Co., 1997. ISBN 0-89791-896-7. doi:[10.1145/258734.258849](https://doi.org/10.1145/258734.258849).
- [71] N. Gelfand, N. J. Mitra, L. J. Guibas, and H. Pottmann. Robust global registration. In *Proceedings of the Symposium on Geometry Processing*, SGP '05. The Eurographics Association, 2005. ISBN 3-905673-24-X. doi:[10.2312/SGP/SGP05/197-206](https://doi.org/10.2312/SGP/SGP05/197-206).
- [72] S. Z. Gilani, A. Mian, and P. Eastwood. Deep, dense and accurate 3D face correspondence for generating population specific deformable models. *Pattern Recognition*, 69:238–250, 2017. ISSN 0031-3203. doi:[10.1016/j.patcog.2017.04.013](https://doi.org/10.1016/j.patcog.2017.04.013).
- [73] T. Groueix. 3D-CODED: 3D correspondences by deep deformation. GitHub, 2018. URL <https://github.com/ThibaultGROUEIX/3D-CODED>.
- [74] T. Groueix, M. Fisher, V. G. Kim, B. C. Russell, and M. Aubry. 3D-CODED: 3D correspondences by deep deformation. In *European Conference on Computer Vision*, ECCV '18. Springer, September 2018. doi:[10.1007/978-3-030-01216-8_15](https://doi.org/10.1007/978-3-030-01216-8_15).
- [75] J. Guo, H. Wang, Z. Cheng, X. Zhang, and D.-M. Yan. Learning local shape descriptors for computing non-rigid dense correspondence. *Computational Visual Media*, 6(1):95–112, 2020. ISSN 2096-0662. doi:[10.1007/s41095-020-0163-y](https://doi.org/10.1007/s41095-020-0163-y).
- [76] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, J. Wan, and N. M. Kwok. A comprehensive performance evaluation of 3D local feature descriptors. *International Journal of Computer Vision*, 116(1):66–89, January 2016. ISSN 1573-1405. doi:[10.1007/s11263-015-0824-y](https://doi.org/10.1007/s11263-015-0824-y).
- [77] O. Halimi, O. Litany, E. R. Rodolà, A. M. Bronstein, and R. Kimmel. Unsupervised learning of dense shape correspondence. In *Conference on Computer Vision and Pattern Recognition*, CVPR '19, pages 4365–4374. IEEE/CVF, 2019. doi:[10.1109/CVPR.2019.00450](https://doi.org/10.1109/CVPR.2019.00450).
- [78] N. Hasler, C. Stoll, M. Sunkel, B. Rosenhahn, and H.-P. Seidel. A statistical model of human pose and body shape. *Computer Graphics Forum*, 28(2):337–346, 2009. doi:[10.1111/j.1467-8659.2009.01373.x](https://doi.org/10.1111/j.1467-8659.2009.01373.x).

- [79] D. L. Hill, D. J. Hawkes, J. E. Crossman, M. J. Gleeson, T. C. Cox, E. E. Bracey, A. J. Strong, and P. Graves. Registration of MR and CT images for skull base surgery using point-like anatomical features. *The British journal of radiology*, 64(767):1030–1035, November 1991. ISSN 0007-1285. doi:[10.1259/0007-1285-64-767-1030](https://doi.org/10.1259/0007-1285-64-767-1030).
- [80] P. W. Holland and R. E. Welsch. Robust regression using iteratively reweighted least-squares. *Communications in Statistics – Theory and Methods*, 6(9):813–827, 1977. doi:[10.1080/03610927708827533](https://doi.org/10.1080/03610927708827533).
- [81] H. Huang, E. Kalogerakis, S. Chaudhuri, D. Ceylan, V. G. Kim, and E. Yumer. Learning local shape descriptors from part correspondences with multiview convolutional networks. *Transactions on Graphics*, 37(1), November 2017. ISSN 0730-0301. doi:[10.1145/3137609](https://doi.org/10.1145/3137609).
- [82] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Computer Graphics Forum*, 27(5):1449–1457, 2008. doi:[10.1111/j.1467-8659.2008.01285.x](https://doi.org/10.1111/j.1467-8659.2008.01285.x).
- [83] Q.-X. Huang, G.-X. Zhang, L. Gao, S.-M. Hu, A. Butscher, and L. Guibas. An optimization approach for extracting and encoding consistent maps in a shape collection. *Transactions on Graphics*, 31(6), November 2012. ISSN 0730-0301. doi:[10.1145/2366145.2366186](https://doi.org/10.1145/2366145.2366186).
- [84] Q.-X. Huang, F. Wang, and L. Guibas. Functional map networks for analyzing and exploring large shape collections. *Transactions on Graphics*, 33(4), July 2014. ISSN 0730-0301. doi:[10.1145/2601097.2601111](https://doi.org/10.1145/2601097.2601111).
- [85] R. Huang, J. Ren, P. Wonka, and M. Ovsjanikov. Consistent ZoomOut: Efficient spectral map synchronization. *Computer Graphics Forum*, 39(5):265–278, 2020. doi:[10.1111/cgf.14084](https://doi.org/10.1111/cgf.14084).
- [86] T. Jiang. *Non-isometric 3D shape registration*. PhD thesis, Media & Communication, Bournemouth University, August 2019.
- [87] T. Jiang, K. Qian, S. Liu, J. Wang, X. Yang, and J. Zhang. Consistent as-similar-as-possible non-isometric surface registration. *The Visual Computer*, 33(6):891–901, 2017. ISSN 1432-2315. doi:[10.1007/s00371-017-1390-9](https://doi.org/10.1007/s00371-017-1390-9).

-
- [88] T. Jiang, X. Yang, J. Zhang, F. Tian, S. Liu, N. Xiang, and K. Qian. Huber- L_1 -based non-isometric surface registration. *The Visual Computer*, 35(6):935–948, May 2019. ISSN 1432-2315. doi:[10.1007/s00371-019-01670-1](https://doi.org/10.1007/s00371-019-01670-1).
- [89] A. E. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3D scenes. *Transactions on Pattern Analysis and Machine Intelligence*, 21(5):433–449, 1999. doi:[10.1109/34.765655](https://doi.org/10.1109/34.765655).
- [90] T. Jost and H. Hügli. Fast ICP algorithms for shape registration. In *DAGM Symposium on Pattern Recognition – Lecture Notes in Computer Science*, volume 2449 of *DAGM '02*, pages 91–99, Berlin, Heidelberg, 2002. Springer-Verlag. ISBN 354044209X. doi:[10.1007/3-540-45783-6_12](https://doi.org/10.1007/3-540-45783-6_12).
- [91] T. Ju, S. Schaefer, and J. Warren. Mean value coordinates for closed triangular meshes. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, pages 561–566. Association for Computing Machinery, 2005. ISBN 9781450378253. doi:[10.1145/1186822.1073229](https://doi.org/10.1145/1186822.1073229).
- [92] P. Kamousi, S. Lazard, A. Maheshwari, and S. Wuhrer. Analysis of farthest point sampling for approximating geodesics in a graph. *Computational Geometry*, 57:1–7, 2016. ISSN 0925-7721. doi:[10.1016/j.comgeo.2016.05.005](https://doi.org/10.1016/j.comgeo.2016.05.005).
- [93] I. Kemelmacher and R. Basri. Indexing with unknown illumination and pose. In *Conference on Computer Vision and Pattern Recognition*, volume 1 of *CVPR '05*, pages 909–916. IEEE, June 2005. doi:[10.1109/CVPR.2005.186](https://doi.org/10.1109/CVPR.2005.186).
- [94] M. Kilian, N. J. Mitra, and H. Pottmann. Geometric modeling in shape space. *Transactions on Graphics*, 26(3), July 2007. ISSN 0730-0301. doi:[10.1145/1276377.1276457](https://doi.org/10.1145/1276377.1276457).
- [95] V. G. Kim, Y. Lipman, and T. Funkhouser. Blended intrinsic maps. *Transactions on Graphics*, 30(4), July 2011. ISSN 0730-0301. doi:[10.1145/2010324.1964974](https://doi.org/10.1145/2010324.1964974).
- [96] V. G. Kim, W. Li, N. J. Mitra, S. DiVerdi, and T. Funkhouser. Exploring collections of 3D models using fuzzy correspondences. *Transactions on Graphics*, 31(4), July 2012. ISSN 0730-0301. doi:[10.1145/2185520.2185550](https://doi.org/10.1145/2185520.2185550).
- [97] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences*, 95(15):8431–8435, 1998. ISSN 0027-8424. doi:[10.1073/pnas.95.15.8431](https://doi.org/10.1073/pnas.95.15.8431).

- [98] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. *Proceedings of the National Academy of Sciences*, 95(15):8431–8435, 1998. ISSN 0027-8424. doi:[10.1073/pnas.95.15.8431](https://doi.org/10.1073/pnas.95.15.8431).
- [99] E. King, X. Y. Xu, A. D. Hughes, Q. Long, S. A. Thom, and K. H. Parker. Quantification of the non-planarity of the human carotid bifurcation. *Biorheology*, 39(3-4):419–424, 2002. ISSN 0006-355X.
- [100] Y. Kleiman and M. Ovsjanikov. Robust structure-based shape correspondence. *Computer Graphics Forum*, 38(1):7–20, April 2019. doi:[10.1111/cgf.13389](https://doi.org/10.1111/cgf.13389).
- [101] A. Kovnatsky, M. M. Bronstein, X. Bresson, and P. Vandergheynst. Functional correspondence by matrix completion. In *Conference on Computer Vision and Pattern Recognition*, CVPR '15, pages 905–914. IEEE, 2015. doi:[10.1109/CVPR.2015.7298692](https://doi.org/10.1109/CVPR.2015.7298692).
- [102] Z. Löhner, E. Rodolà, M. M. Bronstein, D. Cremers, O. Burghard, L. Cosmo, A. Dieckmann, R. Klein, and Y. Sahillioglu. Matching of deformable shapes with topological noise. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '16, pages 55–60. The Eurographics Association, May 2016. ISBN 978-3-03868-004-8. doi:[10.2312/3dor.20161088](https://doi.org/10.2312/3dor.20161088).
- [103] Z. Löhner. Efficient deformable shape correspondence via kernel matching. GitHub, 2017. URL <https://github.com/zorah/KernelMatching>.
- [104] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *International Conference on Computer Vision*, volume 2 of *ICCV '05*, pages 1482–1489 Vol. 2. IEEE/CVF, 2005. doi:[10.1109/ICCV.2005.20](https://doi.org/10.1109/ICCV.2005.20).
- [105] T. Lescoat, H.-T. D. Liu, J.-M. Thiery, A. Jacobson, T. Boubekur, and M. Ovsjanikov. Spectral mesh simplification. *Computer Graphics Forum*, 39(2):315–324, 2020. doi:[10.1111/cgf.13932](https://doi.org/10.1111/cgf.13932).
- [106] Z. Levi and C. Gotsman. Smooth rotation enhanced as-rigid-as-possible mesh animation. *Transactions on Visualization and Computer Graphics*, 21(2):264–277, 2015. doi:[10.1109/TVCG.2014.2359463](https://doi.org/10.1109/TVCG.2014.2359463).
- [107] H. Li, R. W. Sumner, and M. Pauly. Global correspondence optimization for non-rigid registration of depth scans. *Computer Graphics Forum*, 27(5):1421–1430, 2008. doi:[10.1111/j.1467-8659.2008.01282.x](https://doi.org/10.1111/j.1467-8659.2008.01282.x).

- [108] H. Li, B. Adams, L. J. Guibas, and M. Pauly. Robust single-view geometry and motion reconstruction. *Transactions on Graphics*, 28(5):175:1–175:10, December 2009. ISSN 0730-0301. doi:[10.1145/1661412.1618521](https://doi.org/10.1145/1661412.1618521).
- [109] K. Li, J. Yang, Y. K. Lai, and D. Guo. Robust non-rigid registration with reweighted position and transformation sparsity. *Transactions on Visualization and Computer Graphics*, 25(6):2255–2269, June 2018. doi:[10.1109/TVCG.2018.2832136](https://doi.org/10.1109/TVCG.2018.2832136).
- [110] X. Li, C. Wen, L. Wang, and Y. Fang. Topology-constrained shape correspondence. *Transactions on Visualization and Computer Graphics*, pages 1–1, 2020. doi:[10.1109/TVCG.2020.2994013](https://doi.org/10.1109/TVCG.2020.2994013).
- [111] Z. Lian, A. Godil, B. Bustos, M. Daoudi, J. Hermans, S. Kawamura, Y. Kurita, G. Lavoué, H. V. Nguyen, R. Ohbuchi, Y. Ohkita, Y. Ohishi, F. Porikli, M. Reuter, I. Sipiran, D. Smeets, P. Suetens, H. Tabia, and D. Vandermeulen. SHREC '11: Shape retrieval on non-rigid 3D watertight meshes. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '11, pages 79–88. The Eurographics Association, 2011. doi:[10.2312/3DOR/3DOR11/079-088](https://doi.org/10.2312/3DOR/3DOR11/079-088).
- [112] F. A. Limberger and R. C. Wilson. Feature encoding of spectral signatures for 3D non-rigid shape retrieval. In *Proceedings of the British Machine Vision Conference*, BMVC '15, pages 56.1–56.13. BMVA Press, September 2015. ISBN 1-901725-53-7. doi:[10.5244/C.29.56](https://doi.org/10.5244/C.29.56).
- [113] S. Lin, Y. Chen, Y.-K. Lai, R. R. Martin, and Z.-Q. Cheng. Fast capture of textured full-body avatar with RGB-D cameras. *The Visual Computer*, 32(6):681–691, 2016. ISSN 1432-2315. doi:[10.1007/s00371-016-1245-9](https://doi.org/10.1007/s00371-016-1245-9).
- [114] Y. Lipman, D. Levin, and D. Cohen-Or. Green coordinates. *Transactions on Graphics*, 27(3):1–10, August 2008. ISSN 0730-0301. doi:[10.1145/1360612.1360677](https://doi.org/10.1145/1360612.1360677).
- [115] O. Litany, E. Rodolà, A. M. Bronstein, M. M. Bronstein, and D. Cremers. Non-rigid puzzles. *Computer Graphics Forum*, 35(5):135–143, 2016. doi:[10.1111/cgf.12970](https://doi.org/10.1111/cgf.12970).
- [116] O. Litany, T. Remez, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In

- International Conference on Computer Vision, ICCV '17*, pages 5660–5668. IEEE, October 2017. doi:[10.1109/ICCV.2017.603](https://doi.org/10.1109/ICCV.2017.603).
- [117] O. Litany, E. Rodolà, A. M. Bronstein, and M. M. Bronstein. Fully Spectral Partial Shape Matching. *Computer Graphics Forum*, 2017. ISSN 1467-8659. doi:[10.1111/cgf.13123](https://doi.org/10.1111/cgf.13123).
- [118] R. Litman and A. M. Bronstein. Learning spectral descriptors for deformable shape correspondence. *Transactions on Pattern Analysis and Machine Intelligence*, 36(1):171–180, 2014. doi:[10.1109/TPAMI.2013.148](https://doi.org/10.1109/TPAMI.2013.148).
- [119] B. Liu, S. Chen, S.-Q. Xin, Y. He, Z. Liu, and J. Zhao. An optimization-driven approach for computing geodesic paths on triangle meshes. *Computer-Aided Design*, 90:105–112, 2017. ISSN 0010-4485. doi:[10.1016/j.cad.2017.05.022](https://doi.org/10.1016/j.cad.2017.05.022).
- [120] H.-T. D. Liu, A. Jacobson, and M. Ovsjanikov. Spectral coarsening of geometric operators. *Transactions on Graphics*, 38(4), July 2019. ISSN 0730-0301. doi:[10.1145/3306346.3322953](https://doi.org/10.1145/3306346.3322953).
- [121] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: A skinned multi-person linear model. *Transactions on Graphics*, 34(6), October 2015. doi:[10.1145/2816795.2818013](https://doi.org/10.1145/2816795.2818013).
- [122] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, August 1987. ISSN 0097-8930. doi:[10.1145/37402.37422](https://doi.org/10.1145/37402.37422).
- [123] K.-L. Low. Linear least-squares optimization for point-to-plane ICP surface registration. Technical report, TR04-004, Department of Computer Science, University of North Carolina at Chapel Hill, February 2004.
- [124] D. G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 1150–1157, 1999. doi:[10.1109/ICCV.1999.790410](https://doi.org/10.1109/ICCV.1999.790410).
- [125] R. Marin, S. Melzi, E. Rodolà, and U. Castellani. High-resolution augmentation for automatic template-based matching of human models. In *International Conference on 3D Vision, 3DV '19*, pages 230–239. IEEE, September 2019. doi:[10.1109/3DV.2019.00034](https://doi.org/10.1109/3DV.2019.00034).

- [126] R. Marin, S. Melzi, E. Rodolà, and U. Castellani. FARM: Functional automatic registration method for 3D human bodies. *Computer Graphics Forum*, 39(1): 160–173, February 2020. doi:[10.1111/cgf.13751](https://doi.org/10.1111/cgf.13751).
- [127] R. Marin, S. Melzi, P. Musoni, F. Bardoni, M. Tarini, and U. Castellani. CMH: Coordinates manifold harmonics for functional remeshing. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '19, pages 63–70. The Eurographics Association and John Wiley & Sons, Ltd., 2019. doi:[10.2312/3dor.20191063](https://doi.org/10.2312/3dor.20191063).
- [128] J. Masci, D. Boscaini, M. Bronstein, and P. Vandergheynst. ShapeNet: Convolutional neural networks on non-euclidean manifolds. Unpublished, 2015. URL <http://infoscience.epfl.ch/record/204949>.
- [129] J. Masci, D. Boscaini, M. M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *International Conference on Computer Vision Workshops*, ICCVW '15, pages 832–840. IEEE, December 2015. doi:[10.1109/ICCVW.2015.112](https://doi.org/10.1109/ICCVW.2015.112).
- [130] H. Masuda. Topological operators and boolean operations for complex-based non-manifold geometric models. *Computer-Aided Design*, 25(2):119 – 129, 1993. ISSN 0010-4485. doi:[10.1016/0010-4485\(93\)90097-8](https://doi.org/10.1016/0010-4485(93)90097-8).
- [131] S. Melzi, M. Ovsjanikov, G. Roffo, M. Cristani, and U. Castellani. Discrete time evolution process descriptor for shape analysis and matching. *Transactions on Graphics*, 37(1):4:1–4:18, January 2018. ISSN 0730-0301. doi:[10.1145/3144454](https://doi.org/10.1145/3144454).
- [132] S. Melzi, R. Marin, E. Rodolà, U. Castellani, J. Ren, A. Poulénard, P. Wonka, and M. Ovsjanikov. Matching humans with different connectivity. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '19, pages 121–128. The Eurographics Association, May 2019. doi:[10.2312/3dor.20191070](https://doi.org/10.2312/3dor.20191070).
- [133] S. Melzi, J. Ren, E. Rodolà, A. Sharma, P. Wonka, and M. Ovsjanikov. ZoomOut: Spectral upsampling for efficient shape correspondence. *Transactions on Graphics*, 38(6):155:1–155:14, November 2019. ISSN 0730-0301. doi:[10.1145/3355089.3356524](https://doi.org/10.1145/3355089.3356524).
- [134] S. Melzi, R. Marin, P. Musoni, F. Bardoni, M. Tarini, and U. Castellani. Intrinsic/extrinsic embedding for functional remeshing of 3D shapes. *Computers & Graphics*, 88:1–12, 2020. doi:[10.1016/j.cag.2020.02.002](https://doi.org/10.1016/j.cag.2020.02.002).

- [135] J. S. B. Mitchell, D. M. Mount, and C. H. Papadimitriou. The discrete geodesic problem. *SIAM Journal on Computing*, 16:647–668, 1987. doi:[10.1137/0216045](https://doi.org/10.1137/0216045).
- [136] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *Conference on Computer Vision and Pattern Recognition*, CVPR '17, pages 5425–5434. IEEE, 2017. doi:[10.1109/CVPR.2017.576](https://doi.org/10.1109/CVPR.2017.576).
- [137] H. Moon and P. J. Phillips. Computational and performance aspects of PCA-based face-recognition algorithms. *Perception*, 30(3):303–321, 2001. doi:[10.1068/p2896](https://doi.org/10.1068/p2896).
- [138] A. Myronenko and X. Song. Point set registration: Coherent point drift. *Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, 2010. doi:[10.1109/TPAMI.2010.46](https://doi.org/10.1109/TPAMI.2010.46).
- [139] T. Neumann, K. Varanasi, C. Theobalt, M. Magnor, and M. Wacker. Compressed manifold modes for mesh processing. *Computer Graphics Forum*, 33(5):35–44, 2014. doi:[10.1111/cgf.12429](https://doi.org/10.1111/cgf.12429).
- [140] R. A. Newcombe, D. Fox, and S. M. Seitz. DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Conference on Computer Vision and Pattern Recognition*, CVPR '15. IEEE, June 2015. doi:[10.1109/CVPR.2015.7298631](https://doi.org/10.1109/CVPR.2015.7298631).
- [141] A. Nguyen, M. Ben-Chen, K. Welnicka, Y. Ye, and L. Guibas. An optimization approach to improving collections of shape maps. *Computer Graphics Forum*, 30(5):1481–1491, 2011. doi:[10.1111/j.1467-8659.2011.02022.x](https://doi.org/10.1111/j.1467-8659.2011.02022.x).
- [142] D. Nogneng and M. Ovsjanikov. Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum*, 36(2):259–267, May 2017. ISSN 0167-7055. doi:[10.1111/cgf.13124](https://doi.org/10.1111/cgf.13124).
- [143] D. Nogneng, S. Melzi, E. Rodolà, U. Castellani, M. M. Bronstein, and M. Ovsjanikov. Improved functional mappings via product preservation. *Computer Graphics Forum*, 37:179–190, 2018. doi:[10.1111/cgf.13352](https://doi.org/10.1111/cgf.13352).
- [144] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. *Transactions on Graphics*, 31(4):30:1–30:11, July 2012. ISSN 0730-0301. doi:[10.1145/2185520.2185526](https://doi.org/10.1145/2185520.2185526).

- [145] M. Ovsjanikov, Q. Mérigot, V. Pătrăucean, and L. Guibas. Shape matching via quotient spaces. In *Proceedings of the Symposium on Geometry Processing*, SGP '13, pages 1–11, Goslar, DEU, 2013. The Eurographics Association. doi:[10.1111/cgf.12167](https://doi.org/10.1111/cgf.12167).
- [146] M. Ovsjanikov, É. Corman, M. Bronstein, E. Rodolà, M. Ben-Chen, L. Guibas, F. Chazal, and A. Bronstein. Computing and processing correspondences with functional maps. In *ACM SIGGRAPH 2017 Courses*, SIGGRAPH '17. Association for Computing Machinery, 2017. ISBN 9781450350143. doi:[10.1145/3084873.3084877](https://doi.org/10.1145/3084873.3084877).
- [147] C. Papazov and D. Burschka. Deformable 3D shape registration based on local similarity transforms. *Computer Graphics Forum*, 30(5):1493–1502, 2011. doi:[10.1111/j.1467-8659.2011.02023.x](https://doi.org/10.1111/j.1467-8659.2011.02023.x).
- [148] M. Pauly, N. J. Mitra, J. Giesen, M. Gross, and L. J. Guibas. Example-based 3D scan completion. In *Proceedings of the Symposium on Geometry Processing*, SGP '05. The Eurographics Association, 2005. ISBN 3-905673-24-X. doi:[10.2312/SGP/SGP05/023-032](https://doi.org/10.2312/SGP/SGP05/023-032).
- [149] D. Pickup, X. Sun, P. L. Rosin, R. R. Martin, Z. Cheng, Z. Lian, M. Aono, A. B. Hamza, A. M. Bronstein, M. M. Bronstein, S. Bu, U. Castellani, S. Cheng, V. Garro, A. Giachetti, A. Godil, J. Han, H. Johan, L. Lai, B. Li, C. Li, H. Li, R. Litman, X. Liu, Z. Liu, Y. Lu, A. Tatsuma, and J. Ye. Shape retrieval of non-rigid 3D human models. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '14, pages 101–110. The Eurographics Association, 2014. ISBN 978-3-905674-58-3. doi:[10.2312/3dor.20141056](https://doi.org/10.2312/3dor.20141056).
- [150] M. Pinheiro, X. Ma, M. J. Fagan, G. T. McIntyre, P. Lin, G. Sivamurthy, and P. A. Mossey. A 3D cephalometric protocol for the accurate quantification of the craniofacial symmetry and facial growth. *Journal of Biological Engineering*, 13(1):42, 2019. ISSN 1754-1611. doi:[10.1186/s13036-019-0171-6](https://doi.org/10.1186/s13036-019-0171-6).
- [151] L. Pishchulin, S. Wuhrer, T. Helten, C. Theobalt, and B. Schiele. Building statistical shape spaces for 3D human modeling. *Pattern Recognition*, 67:276–286, 2017. ISSN 0031-3203. doi:[10.1016/j.patcog.2017.02.018](https://doi.org/10.1016/j.patcog.2017.02.018).
- [152] J. Pokrass, A. M. Bronstein, and M. M. Bronstein. Partial shape matching without point-wise correspondence. *Numerical Mathematics: Theory, Methods and Applications*, 6(1):223–244, 2013. doi:[10.1017/S1004897900001203](https://doi.org/10.1017/S1004897900001203).

- [153] G. Pons-Moll, J. Romero, N. Mahmood, and M. J. Black. Dyna: A model of dynamic human shape in motion. *Transactions on Graphics*, 34(4), July 2015. ISSN 0730-0301. doi:[10.1145/2766993](https://doi.org/10.1145/2766993).
- [154] T. Popa, D. Julius, and A. Sheffer. Material-aware mesh deformations. In *International Conference on Shape Modeling and Applications*, SMI '06, pages 22–22. IEEE, June 2006. doi:[10.1109/SMI.2006.23](https://doi.org/10.1109/SMI.2006.23).
- [155] J. Ren, A. Poulénard, P. Wonka, and M. Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *Transactions on Graphics*, 37:248:1–248:16, December 2018. doi:[10.1145/3272127.3275040](https://doi.org/10.1145/3272127.3275040).
- [156] J. Ren, M. Panine, P. Wonka, and M. Ovsjanikov. Structured regularization of functional map computations. *Computer Graphics Forum*, 38(5):39–53, 2019. doi:[10.1111/cgf.13788](https://doi.org/10.1111/cgf.13788).
- [157] K. M. Robinette, H. Daanen, and E. Paquet. The CAESAR project: A 3D surface anthropometry survey. In *International Conference on 3D Digital Imaging and Modeling*, 3DIM '99, pages 380–386. IEEE, August 1999. ISBN 0769500625. doi:[10.1109/IM.1999.805368](https://doi.org/10.1109/IM.1999.805368).
- [158] E. Rodolà, S. Bulò, T. Windheuser, M. Vestner, and D. Cremers. Dense non-rigid shape correspondence using random forests. In *Conference on Computer Vision and Pattern Recognition*, CVPR '14, pages 4177–4184. IEEE, 2014. doi:[10.1109/CVPR.2014.532](https://doi.org/10.1109/CVPR.2014.532).
- [159] E. Rodolà, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. *Computer Graphics Forum*, 36(1):222–236, January 2017. ISSN 0167-7055. doi:[10.1111/cgf.12797](https://doi.org/10.1111/cgf.12797).
- [160] E. Rodolà, L. Cosmo, O. Litany, M. M. Bronstein, A. M. Bronstein, N. Audebert, A. B. Hamza, A. Boulch, U. Castellani, M. N. Do, A.-D. Duong, T. Furuya, A. Gasparetto, Y. Hong, J. Kim, B. L. Saux, R. Litman, M. Masoumi, G. Minello, H.-D. Nguyen, V.-T. Nguyen, R. Ohbuchi, V.-K. Pham, T. V. Phan, M. Rezaei, A. Torsello, M.-T. Tran, Q.-T. Tran, B. Truong, L. Wan, and C. Zou. Deformable shape retrieval with missing parts. In *Eurographics Workshop on 3D Object Retrieval*, 3DOR '17, pages 85–94. The Eurographics Association, April 2017. ISBN 978-3-03868-030-7. doi:[10.2312/3dor.20171057](https://doi.org/10.2312/3dor.20171057).

- [161] E. Rodolà, M. Moeller, and D. Cremers. Regularized pointwise map recovery from functional correspondence. *Computer Graphics Forum*, 36(8):700–711, 2017. doi:[10.1111/cgf.13160](https://doi.org/10.1111/cgf.13160).
- [162] E. Rodolà, M. Moeller, and D. Cremers. Point-wise map recovery and refinement from functional correspondence. In *Vision, Modeling & Visualization*, pages 25–32. The Eurographics Association, 2015. ISBN 978-3-905674-95-8. doi:[10.2312/vmv.20151254](https://doi.org/10.2312/vmv.20151254).
- [163] R. Rostami, F. S. Bashiri, B. Rostami, and Z. Yu. A survey on data-driven 3D shape descriptors. *Computer Graphics Forum*, 38(1):356–393, 2019. doi:[10.1111/cgf.13536](https://doi.org/10.1111/cgf.13536).
- [164] J. Roufousse, A. Sharma, and M. Ovsjanikov. Unsupervised deep learning for structured shape matching. In *International Conference on Computer Vision, ICCV '19*, pages 1617–1627. IEEE/CVF, 2019.
- [165] S. Rusinkiewicz and M. Levoy. Efficient variants of the ICP algorithm. In *International Conference on 3D Digital Imaging and Modeling, 3DIM '01*, pages 145–152. IEEE, 2001. doi:[10.1109/IM.2001.924423](https://doi.org/10.1109/IM.2001.924423).
- [166] Russian3DScanner. Wrap 3.3, April 2018. URL <https://www.russian3dscanner.com/>.
- [167] Russian3DScanner. Wrap 3.4, June 2019. URL <https://www.russian3dscanner.com/>.
- [168] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz. Aligning point cloud views using persistent feature histograms. In *RSJ International Conference on Intelligent Robots and Systems, IROS '08*, pages 3384–3391. IEEE, 2008. doi:[10.1109/IROS.2008.4650967](https://doi.org/10.1109/IROS.2008.4650967).
- [169] R. B. Rusu, N. Blodow, and M. Beetz. Fast Point Feature Histograms (FPFH) for 3D registration. In *International Conference on Robotics and Automation, ICRA '09*, pages 3212–3217. IEEE, 2009. doi:[10.1109/ROBOT.2009.5152473](https://doi.org/10.1109/ROBOT.2009.5152473).
- [170] Y. Sahillioğlu and Y. Yemez. Partial 3D correspondence from shape extremities. *Computer Graphics Forum*, 33(6):63–76, 2014. doi:[10.1111/cgf.12278](https://doi.org/10.1111/cgf.12278).
- [171] Y. Sahillioğlu. Recent advances in shape correspondence. *The Visual Computer*, September 2019. ISSN 1432-2315. doi:[10.1007/s00371-019-01760-0](https://doi.org/10.1007/s00371-019-01760-0).

- [172] Y. Sahillioğlu. A genetic isometric shape correspondence algorithm with adaptive sampling. *Transactions on Graphics*, 37(5):175:1–175:14, 2018. doi:[10.1145/3243593](https://doi.org/10.1145/3243593).
- [173] Y. Sahillioğlu and Y. Yemez. Scale normalization for isometric shape matching. *Computer Graphics Forum*, 31(7):2233–2240, 2012. doi:[10.1111/j.1467-8659.2012.03216.x](https://doi.org/10.1111/j.1467-8659.2012.03216.x).
- [174] Y. Sahillioğlu and L. Kavan. Skuller: A volumetric shape registration algorithm for modeling skull deformities. *Medical image analysis*, 23(1):15–27, July 2015. ISSN 1361-8423. doi:[10.1016/j.media.2015.03.005](https://doi.org/10.1016/j.media.2015.03.005).
- [175] A. Salazar, S. Wuhler, C. Shu, and F. Prieto. Fully automatic expression-invariant face correspondence. *Machine Vision and Applications*, 25(4):859–879, 2014. ISSN 14321769. doi:[10.1007/s00138-013-0579-9](https://doi.org/10.1007/s00138-013-0579-9).
- [176] P. Sand, L. McMillan, and J. Popović. Continuous capture of skin deformation. *Transactions on Graphics*, 22(3):578–586, July 2003. doi:[10.1145/882262.882310](https://doi.org/10.1145/882262.882310).
- [177] B. Sauvage, S. Hahmann, G.-P. Bonneau, and G. Elber. Detail preserving deformation of B-spline surfaces with volume constraint. *Computer Aided Geometric Design*, 25(8):678–696, November 2008. ISSN 0167-8396. doi:[10.1016/j.cagd.2007.10.004](https://doi.org/10.1016/j.cagd.2007.10.004).
- [178] H. Schlichtkrull. Curves and surfaces. University of Copenhagen, 2013. URL <https://web.archive.org/web/20201224212133/http://web.math.ku.dk/noter/filer/geom1.pdf>.
- [179] F. Schmidt, F. Phillips, and R. W. Fleming. Visual perception of shape-transforming processes: ‘shape scission’. *Cognition*, 189:167–180, August 2019. ISSN 0010-0277. doi:[10.1016/j.cognition.2019.04.006](https://doi.org/10.1016/j.cognition.2019.04.006).
- [180] J. A. Schnabel, D. Rueckert, M. Quist, J. M. Blackall, A. D. Castellano-Smith, T. Hartkens, G. P. Penney, W. A. Hall, H. Liu, C. L. Truwit, F. A. Gerritsen, D. L. G. Hill, and D. J. Hawkes. A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations. In *Medical Image Computing and Computer-Assisted Intervention*, MICCAI ’01, pages 573–581, Berlin, Heidelberg, 2001. Springer. ISBN 978-3-540-45468-7. doi:[10.1007/3-540-45468-3_69](https://doi.org/10.1007/3-540-45468-3_69).

- [181] M. Shoham, A. Vaxman, and M. Ben-Chen. Hierarchical functional maps between subdivision surfaces. *Computer Graphics Forum*, 38(5):55–73, 2019. doi:[10.1111/cgf.13789](https://doi.org/10.1111/cgf.13789).
- [182] K. Siddiqi, J. Zhang, D. Macrini, A. Shokoufandeh, S. Bouix, and S. Dickinson. Retrieving articulated 3D models using medial surfaces. *Machine Vision and Applications*, 19(4):261–275, May 2008. doi:[10.1007/s00138-007-0097-8](https://doi.org/10.1007/s00138-007-0097-8).
- [183] C. Sinthanayothin and W. Tharanon. 3D-3D registration: Surface rendering plus skull and soft tissue registration. In *Conference on Industrial Electronics and Applications*, ICIEA '06, pages 1–5. IEEE, 2006. doi:[10.1109/ICIEA.2006.257340](https://doi.org/10.1109/ICIEA.2006.257340).
- [184] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.-P. Seidel. Laplacian surface editing. In *Proceedings of the Symposium on Geometry Processing*, SGP '04. The Eurographics Association, 2004. ISBN 3-905673-13-4. doi:[10.2312/SGP/SGP04/179-188](https://doi.org/10.2312/SGP/SGP04/179-188).
- [185] O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Proceedings of the Symposium on Geometry Processing*, SGP '07, pages 109–116. The Eurographics Association and John Wiley & Sons, Ltd., 2007. doi:[10.2312/SGP/SGP07/109-116](https://doi.org/10.2312/SGP/SGP07/109-116).
- [186] C. Stride. Creating a dataset of 3D deformed shapes for evaluating 3D model registration techniques. Cardiff University, May 2019. Bachelor’s thesis.
- [187] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. *Transactions on Graphics*, 23(3):399–405, August 2004. doi:[10.1145/1015706.1015736](https://doi.org/10.1145/1015706.1015736).
- [188] R. W. Sumner, J. Schmid, and M. Pauly. Embedded deformation for shape manipulation. *Transactions on Graphics*, 26(3):80–es, July 2007. ISSN 0730-0301. doi:[10.1145/1276377.1276478](https://doi.org/10.1145/1276377.1276478).
- [189] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28:1383–1392, August 2009. doi:[10.1111/j.1467-8659.2009.01515.x](https://doi.org/10.1111/j.1467-8659.2009.01515.x).
- [190] Z. Sun, Y. He, A. Gritsenko, A. Lendasse, and S. Baek. Embedded spectral descriptors: learning the point-wise correspondence metric via Siamese neural

- networks. *Journal of Computational Design and Engineering*, 7(1):18–29, March 2020. ISSN 2288-5048. doi:[10.1093/jcde/qwaa003](https://doi.org/10.1093/jcde/qwaa003).
- [191] G. K. L. Tam, Z. Cheng, Y.-K. Lai, F. C. Langbein, Y. Liu, D. Marshall, R. R. Martin, X. Sun, and P. L. Rosin. Registration of 3D point clouds and meshes: A survey from rigid to nonrigid. *Transactions on Visualization and Computer Graphics*, 19(7):1199–1217, July 2013. ISSN 1077-2626. doi:[10.1109/TVCG.2012.310](https://doi.org/10.1109/TVCG.2012.310).
- [192] G. K. L. Tam, R. R. Martin, P. L. Rosin, and Y.-K. Lai. Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry. *Transactions on Graphics*, 33(1):4:1–4:17, February 2014. ISSN 0730-0301. doi:[10.1145/2517967](https://doi.org/10.1145/2517967).
- [193] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H. Seidel. Isometric registration of ambiguous and partial data. In *Conference on Computer Vision and Pattern Recognition*, CVPR '09, pages 1185–1192. IEEE, 2009. doi:[10.1109/CVPR.2009.5206775](https://doi.org/10.1109/CVPR.2009.5206775).
- [194] F. Tombari, S. Salti, and L. Di Stefano. Unique signatures of histograms for local surface description. In *European Conference on Computer Vision*, ECCV '10, pages 356–369. Springer, September 2010. ISBN 978-3-642-15558-1. doi:[10.1007/978-3-642-15558-1_26](https://doi.org/10.1007/978-3-642-15558-1_26).
- [195] N. Utkualp and I. Ercan. Anthropometric measurements usage in medical sciences. *BioMed Research International*, 2015:404261:1–404261:7, 2015. ISSN 2314-6133. doi:[10.1155/2015/404261](https://doi.org/10.1155/2015/404261).
- [196] O. van Kaick, H. Zhang, G. Hamarneh, and D. Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, July 2011. doi:[10.1111/j.1467-8659.2011.01884.x](https://doi.org/10.1111/j.1467-8659.2011.01884.x).
- [197] O. van Kaick, H. Zhang, and G. Hamarneh. Bilateral maps for partial matching. *Computer Graphics Forum*, 32(6):189–200, 2013. doi:[10.1111/cgf.12084](https://doi.org/10.1111/cgf.12084).
- [198] D. Vandermeulen, P. Claes, D. Loeckx, S. De Greef, G. Willems, and P. Suetens. Computerized craniofacial reconstruction using CT-derived implicit surface representations. *Forensic science international*, 159 Suppl 1:S164–74, May 2006. ISSN 0379-0738. doi:[10.1016/j.forsciint.2006.02.036](https://doi.org/10.1016/j.forsciint.2006.02.036).

- [199] N. Verma, E. Boyer, and J. Verbeek. FeaStNet: Feature-steered graph convolutions for 3D shape analysis. In *Conference on Computer Vision and Pattern Recognition*, CVF '18, pages 2598–2606. IEEE, 2018. doi:[10.1109/CVPR.2018.00275](https://doi.org/10.1109/CVPR.2018.00275).
- [200] M. Vestner, Z. Löhner, A. Boyarski, O. Litany, R. Slossberg, T. Remez, E. Rodolà, A. M. Bronstein, M. M. Bronstein, R. Kimmel, and D. Cremers. Efficient deformable shape correspondence via kernel matching. In *International Conference on 3D Vision*, 3DV '17, pages 517–526. IEEE, October 2017. doi:[10.1109/3DV.2017.00065](https://doi.org/10.1109/3DV.2017.00065).
- [201] M. Vestner, R. Litman, E. Rodolà, A. M. Bronstein, and D. Cremers. Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Conference on Computer Vision and Pattern Recognition*, CVPR '17, pages 6681–6690. IEEE, July 2017. doi:[10.1109/CVPR.2017.707](https://doi.org/10.1109/CVPR.2017.707).
- [202] D. Vlastic, I. Baran, W. Matusik, and J. Popović. Articulated mesh animation from multi-view silhouettes. *Transactions on Graphics*, 27(3):97:1–97:9, August 2008. ISSN 0730-0301. doi:[10.1145/1399504.1360696](https://doi.org/10.1145/1399504.1360696).
- [203] L. Wei, Q.-X. Huang, D. Ceylan, É. Vouga, and H. Li. Dense human body correspondence using convolutional networks. In *Conference on Computer Vision and Pattern Recognition*, CVPR '16, pages 1544–1553. IEEE, 2016. doi:[10.1109/CVPR.2016.171](https://doi.org/10.1109/CVPR.2016.171).
- [204] Z. Wood, H. Hoppe, M. Desbrun, and P. Schröder. Removing excess topology from isosurfaces. *Transactions on Graphics*, 23(2):190–208, April 2004. ISSN 0730-0301. doi:[10.1145/990002.990007](https://doi.org/10.1145/990002.990007).
- [205] Y. Wu, J. Yang, and J. Zhao. Partial 3D shape functional correspondence via fully spectral eigenvalue alignment and upsampling refinement. *Computers & Graphics*, 2020. ISSN 0097-8493. doi:[10.1016/j.cag.2020.09.004](https://doi.org/10.1016/j.cag.2020.09.004).
- [206] Z. Wu, K. Li, Y. Lai, and J. Yang. Global as-conformal-as-possible non-rigid registration of multi-view scans. In *International Conference on Multimedia and Expo*, ICME '19, pages 308–313. IEEE, 2019. doi:[10.1109/ICME.2019.00061](https://doi.org/10.1109/ICME.2019.00061).
- [207] S. Wuhrer, C. Shu, and P. Xi. Landmark-free posture invariant human shape correspondence. *The Visual Computer*, 27(9):843–852, 2011. ISSN 1432-2315. doi:[10.1007/s00371-011-0557-z](https://doi.org/10.1007/s00371-011-0557-z).

- [208] S. Yamazaki, S. Kagami, and M. Mochimaru. Non-rigid shape registration using similarity-invariant differential coordinates. In *International Conference on 3D Vision*, 3DV '13, pages 191–198. IEEE, 2013. doi:[10.1109/3DV.2013.33](https://doi.org/10.1109/3DV.2013.33).
- [209] J. Yang, Y. Xiao, and Z. Cao. Toward the repeatability and robustness of the local reference frame for 3D shape matching: An evaluation. *Transactions on Image Processing*, 27(8):3766–3781, 2018. doi:[10.1109/TIP.2018.2827330](https://doi.org/10.1109/TIP.2018.2827330).
- [210] J. Yang, K. Li, K. Li, and Y.-K. Lai. Sparse non-rigid registration of 3D shapes. *Computer Graphics Forum*, 34(5):89–99, 2015. ISSN 1467-8659. doi:[10.1111/cgf.12699](https://doi.org/10.1111/cgf.12699).
- [211] Y. Yang, Y. Yu, Y. Zhou, S. Du, J. Davis, and R. Yang. Semantic parametric reshaping of human body models. In *International Conference on 3D Vision*, volume 2 of *3DV '14*, pages 41–48. IEEE, 2014. doi:[10.1109/3DV.2014.47](https://doi.org/10.1109/3DV.2014.47).
- [212] Y. Yao, B. Deng, W. Xu, and J. Zhang. Quasi-newton solver for robust non-rigid registration. In *Conference on Computer Vision and Pattern Recognition*, CVPR '20, pages 7597–7606. IEEE/CVF, 2020. doi:[10.1109/CVPR42600.2020.00762](https://doi.org/10.1109/CVPR42600.2020.00762).
- [213] I.-C. Yeh, C.-H. Lin, O. Sorkine, and T.-Y. Lee. Template-based 3D model fitting using dual-domain relaxation. *Transactions on Visualization and Computer Graphics*, 17(8):1178–1190, August 2011. ISSN 1077-2626. doi:[10.1109/TVCG.2010.124](https://doi.org/10.1109/TVCG.2010.124).
- [214] Y. Yoshiyasu, W.-C. Ma, E. Yoshida, and F. Kanehiro. As-conformal-as-possible surface registration. In *Proceedings of the Symposium on Geometry Processing*, SGP '14, pages 257–267. The Eurographics Association, 2014. doi:[10.1111/cgf.12451](https://doi.org/10.1111/cgf.12451).
- [215] Y. Yoshiyasu, E. Yoshida, and L. Guibas. Symmetry aware embedding for shape correspondence. *Computers & Graphics*, 60:9–22, 2016. ISSN 0097-8493. doi:[10.1016/j.cag.2016.07.002](https://doi.org/10.1016/j.cag.2016.07.002).
- [216] A. Zeng, S. Song, M. Niessner, M. Fisher, J. Xiao, and T. Funkhouser. 3DMatch: Learning the matching of local 3D geometry in range scans. In *Conference on Computer Vision and Pattern Recognition*, CVPR '17, pages 1–12. IEEE, 2017.
- [217] J. Zhang, Y. Yao, and B. Deng. Fast and robust iterative closest point, 2020.

- [218] K. Zhang, Y. Cheng, and W. K. Leow. Dense correspondence of skull models by automatic detection of anatomical landmarks. In *Proceedings, Part I, of the 15th International Conference on Computer Analysis of Images and Patterns*, volume 8047 of *CAIP 2013*, pages 229–236. Springer-Verlag, 2013. ISBN 9783642402609. doi:[10.1007/978-3-642-40261-6_27](https://doi.org/10.1007/978-3-642-40261-6_27).
- [219] J. Zhao, X. Qi, C. Wen, N. Lei, and X. Gu. Automatic and robust skull registration based on discrete uniformization. In *International Conference on Computer Vision, ICCV '19*, pages 431–440. IEEE/CVF, 2019. doi:[10.1109/ICCV.2019.00052](https://doi.org/10.1109/ICCV.2019.00052).
- [220] H. Zhu, B. Guo, K. Zou, Y. Li, K.-V. Yuen, L. Mihaylova, and H. Leung. A review of point set registration: From pairwise registration to groupwise registration. *Sensors*, 19(5):1191, March 2019. ISSN 1424-8220. doi:[10.3390/s19051191](https://doi.org/10.3390/s19051191).
- [221] J. Zhu, Y. Hosaka, and H. Yoshioka. A robust algorithm to remove the self-intersection of 3D mesh data without changing the original shape. *Journal of Physics: Conference Series*, 1314:012149, October 2019. doi:[10.1088/1742-6596/1314/1/012149](https://doi.org/10.1088/1742-6596/1314/1/012149).
- [222] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *Conference on Computer Vision and Pattern Recognition, CVPR '12*, pages 2879–2886. IEEE, 2012. doi:[10.1109/CVPR.2012.6248014](https://doi.org/10.1109/CVPR.2012.6248014).
- [223] M. Zollhöfer, M. Nießner, S. Izadi, C. Rehmann, C. Zach, M. Fisher, C. Wu, A. Fitzgibbon, C. Loop, C. Theobalt, and M. Stamminger. Real-time non-rigid reconstruction using an RGB-D camera. *Transactions on Graphics*, 33(4), July 2014. ISSN 0730-0301. doi:[10.1145/2601097.2601165](https://doi.org/10.1145/2601097.2601165).
- [224] S. Zuffi, A. Kanazawa, D. W. Jacobs, and M. J. Black. 3D Menagerie: Modeling the 3D shape and pose of animals. In *Conference on Computer Vision and Pattern Recognition, CVPR '17*, pages 5524–5532. IEEE, July 2017. doi:[10.1109/cvpr.2017.586](https://doi.org/10.1109/cvpr.2017.586).