

A Soft-Correspondence Approach to Shape-based Disease Grading with Graph Convolutional Networks

Julius Mayer

Daniel Baum

Zuse Institute Berlin

Felix Ambellan

Christoph von Tycowicz

Freie Universität Berlin, Zuse Institute Berlin

MAYER@ZIB.DE

BAUM@ZIB.DE

AMBELLAN@ZIB.DE

VONTYCOWICZ@ZIB.DE

Editor: Editor's name

Abstract

Shape analysis provides principled means for understanding anatomical structures from medical images. The underlying notions of shape spaces, however, come with strict assumptions prohibiting the analysis of incomplete and/or topologically varying shapes. This work aims to alleviate these limitations by adapting the concept of *soft correspondences*. In particular, we present a graph-based learning approach for morphometric classification of disease states that is based on a generalized notion of shape correspondences in terms of functional maps. We demonstrate the performance of the derived classifier on the open-access ADNI database for differentiating normal controls and subjects with Alzheimer's disease. Notably, our experiment shows that our approach can improve over state-of-the-art from geometric deep learning.

Keywords: Shape Analysis, Classification, Functional Maps, Geometric Deep Learning

1. Introduction

Shapes of anatomies and variations thereof pose a key source for the understanding of medical phenomena including physiological processes like growth as well as pathological conditions that are associated to tissue malformations. Concerning the latter, shape-informed reasoning plays a significant role in medical decision making including therapy planning, personalized assessment and stratification for clinical interventions. From a mathematical point of view, shapes are an instance of geometric data that require dedicated computational treatment. In particular, there is increasing evidence that data-analytical tools that account for the inherent geometric structure yield improved consistency and performance. Consequently, there is a strong impetus to generalize established approaches that have been derived for Euclidean data to geometric ones. This concerns deep learning and statistics as well as data processing and visualization alike.

Central to morphological analysis is the comparison of related forms. This requires a coordinization of shapes leading to a notion of shape space in which each point represents a specific shape (Ambellan et al., 2019). An established approach is to consider transformations connecting shapes: An object class under study can be represented by a common deformable template that accounts for the typicality of the objects' structure. The shape variability is then represented by deformations that are applied to the template. In this line of work, approaches based on Riemannian methods have shown promising results for

tasks such as discovery of biomarkers (Tack et al., 2021), risk assessment of clinical outcomes (Ambellan et al., 2021a), and longitudinal analysis (Gerig et al., 2016; Nava-Yazdani et al., 2022).

Despite these advances, frameworks for geometric morphometry still rely on point-to-point correspondences between shapes: Either explicitly in form of homologous landmarks (von Tycowicz et al., 2018) or implicitly in terms of diffeomorphisms of the ambient space (Bauer et al., 2014). Point-to-point correspondences have fundamental limitations that impede the analysis of shape collections with incomplete or topologically varying objects. This is a major problem for the analysis of empirically given sets of shapes (see Fig. 2 for an illustration), since they often contain topological variations (“real” ones as well as those caused by acquisition artifacts or reconstruction errors) or are incomplete (e.g., due to spatial limitations in tomographic reconstructions or due to destruction and decay).

Consequently there is a high interest in developing novel concepts that pose less strict assumptions. While recent work on non-rigid registration allows for partial matching (Antonsanti et al., 2021) and topology changes via user-specified discontinuities (Nielsen et al., 2019), extensions to group-wise analysis are still at an early stage of research. A promising, alternative approach (Ovsjanikov et al., 2012) that we evaluate in this work is to generalize the notion of correspondence between shapes in terms of maps between real-valued functions on the surfaces instead of points thereon. Remarkably, such *functional maps* facilitate simple and efficient shape matching (Melzi et al., 2019) as well as shape difference operators that characterize distortion between shapes (Corman et al., 2017). Furthermore, recent advances for improving the cycle consistency in functional map networks (Wang et al., 2013; Huang et al., 2019) via latent representations give rise to novel notions of soft-correspondence-based shape spaces, the full potential of which remains to be explored.

Another alternative approach to explicit shape spaces is to infer the underlying structure entirely from the data at hand. In this context, deep learning has led to qualitative breakthroughs for various tasks. Still, as shapes are described by curved surfaces, they are geometric objects in their own right and require dedicated neural network units. The study of such units falls into the category of geometric deep learning and we refer to (Bronstein et al., 2017) for an overview.

Aside from the domain of application, learning methods can be categorized into inductive and transductive approaches. Whereas inductive learning tries to infer a general model from labeled examples in order to predict labels of unseen ones, transductive approaches infer labels simultaneously on training and test data and, thus, can exploit patterns in both during the learning phase. Conceptually, transductive learning faces a simpler problem as it avoids solving a more general one as an intermediate step. Recently, graph convolutional networks have been shown to provide effective transductive learning schemes for disease classification from imaging (Parisot et al., 2018) as well as shape-based features (von Tycowicz, 2020).

In this work, we derive a novel, shape-based classification approach that is based on a flexible, yet descriptive notion of shape space and, on the other hand, casts the grading task as semi-supervised node classification problem on a shape-valued graph. In particular, our approach is tolerant to topological variations and incompleteness in shape collections. The proposed framework draws upon geometric deep learning to incorporate the geometric character of shapes as well as to address the irregular structure in sampling patterns found in collections of clinical observations. We evaluate the performance of our model w.r.t.

shape-based classification of hippocampus malformations due to Alzheimer’s disease. In particular, we achieve state-of-the-art accuracies outperforming recent approaches based on functional maps (Huang et al., 2019) as well as geometric deep learning (han).

2. Method

2.1. Functional correspondence

A functional correspondence between two shapes M and N is given by a map which relates real-valued functions on one shape to real-valued functions on the other shape and is therefore called functional map. Let U_M (resp. U_N) be a finite subspace of the real-valued functions on M (resp. N) that is stable under small shape deformations and contains good approximations of smooth functions. Given a basis for U_M, U_N , the functional map from M to N can approximately be encoded by a matrix $C \in \mathbb{R}^{n \times m}$ s.t. $b \approx Ca$ where a, b represent two corresponding functions as coordinates w.r.t. their respective basis.

We compute C by using ZOOMOUT refinement as proposed in (Melzi et al., 2019). As initialization, we take a 10×10 matrix that accounts for correspondences of approximate landmarks obtained via non-rigid registration.

Shape differences. The area-based shape difference V and the conformal-based shape difference R defined by Rustamov et al. (2013) are operators respectively characterizing area and conformal distortion from M to N by altering real-valued functions on M which are supported in a region of the distortion. In the case of $S \in \{M, N\}$ being discrete, let $H_{S,V}$ be the matrix encoding the inner product $(f_1, f_2) \mapsto \int f_1 f_2 d\mu$ and $H_{S,R}$ the matrix encoding the inner product $(f_1, f_2) \mapsto \int (\nabla f_1)^T \nabla f_2 d\mu$ where $f_1, f_2 \in U_S$ and μ is some surface area measure. Then V is encoded by $D_V := H_{M,V}^\dagger C^T H_{N,V} C$ and R is encoded by $D_R := H_{M,R}^\dagger C^T H_{N,R} C$, where \dagger indicates the pseudo-inverse.

The two shape difference operators V and R depend only on the metric (i.e. *first fundamental form* I) and, thus, are determined exclusively in terms of distances on N . A classical result of surface theory (do Carmo, 1976, p. 236), however, shows that a full characterization of the local geometry of a surface also requires knowledge of the differential $D\nu$ of the




Figure 1: Comparison of distance isolines. The standard metric (left) does not take extrinsic geometry into account, in contrast to the isophotic metric (right, with $\alpha = 8$).

surface normals ν . Therefore, we propose a novel extension of the original differences in order to also capture variation in extrinsic geometry. Let $\mathbb{I}(\cdot, \cdot) := \mathbb{I}(D\nu(\cdot), D\nu(\cdot))$ denote the *third fundamental form*. Then, we equip N with the isophotic metric (Pottmann et al., 2004) $\mathbb{I} + \alpha\mathbb{I}$ for which surface distances also depend on the variation of the surface normals along connecting geodesics (see Figure 1). This approach does not require the construction of an offset surface as in (Corman et al., 2017) and additionally provides a commensuration parameter $\alpha \geq 0$ that allows to weight the influence of normal variation.

2.2. Latent functional representation

In order to perform group-wise analysis, we require a consistent description of structural changes at the population level, i.e. shape differences need to be expressed w.r.t. a common reference frame. To this end, we employ concepts from map synchronization that allow us to obtain such a reference system in terms of latent spaces. Throughout this paper, we assume to be given a shape collection $\mathcal{S} = \{S_1, \dots, S_n\}$ represented as triangle meshes. We choose U_{S_i} to be the space spanned by the first 50 eigenfunctions of the Laplace-Beltrami operator discretized as in (Botsch et al., 2010) using the standard weighted cotangent scheme. Let Λ_i denote the diagonal matrix of its first m eigenvalues in rising order and Φ_i denote the matrix containing the respective eigenfunctions as columns.

Functional map network. A functional map network (FMN) of the shape collection \mathcal{S} is a connected directed graph with the shapes as nodes and edges E between shapes that are sufficiently similar.

We choose the Euclidean distance between the shape-DNA descriptors (Reuter et al., 2006) as a measure for similarity and connect each shape with the k nearest shapes, where k is large enough to guarantee the FMN to be connected. Since the computation of the latent bases is simplified if the FMN is undirected, we add the edge (j, i) to every $(i, j) \in E$ and set all weights of existing edges to the value 1. Note, the FMN is not used as graph for the later applied graph neural network, but only to compute the latent bases.

Latent basis. Let C_{ij} be the functional correspondence matrix from S_i to S_j for $(i, j) \in E$. Commonalities among a shape collection with a given FMN are represented by families of functions that are consistent throughout the whole collection, i.e. $y_j = C_{ij}y_i$ for all $(i, j) \in E$ where y_1, \dots, y_n are coordinates of functions of such a family. A basis for the space of all functions in U_i that belong to such a family is called latent basis. Let Y_i denote a matrix such that its columns are coordinates of functions forming a latent basis on S_i .

Following (Wang et al., 2013) and (Huang et al., 2019), we compute Y_i by minimizing $\sum_{(i,j) \in E} \|C_{ij}Y_i - Y_j\|$ s.t. $\sum_{i=1}^n Y_i^T Y_i = I$ and $\sum_{i=1}^n Y_i^T \Lambda_i Y_i$ is a diagonal matrix. This minimization problem defines a basis called canonical consistent latent basis (CCLB). Since commonly less than 50 functions exist per latent basis, we restrict Y_i to 40 columns only.

Limit shape. As examined in (Huang et al., 2019), the CCLB of S_i merged to Y_i can be seen as a functional correspondence matrix from an only implicitly given shape called limit shape to S_i . The eigenvalues of $\sum_{i=1}^n Y_i^T \Lambda_i Y_i$ are called the spectrum of the limit shape. Let Λ_0 denote the diagonal matrix of these eigenvalues in rising order.

Eventually, based on the geometry of the limit shape, we can express the shape differences in this common frame by

$$D_i^V = Y_i^T Y_i \quad \text{and} \quad (1)$$

$$D_i^R = \Lambda_0^\dagger Y_i^T \Lambda_i Y_i. \quad (2)$$

2.3. Transductive learning

In the following, we derive a neural network model that is conditioned on an entire collection of shapes, each of which belongs to an individual for which we want to predict a disease state. Typically, the sampling in clinical data sets does not follow a regular grid and there exists no natural ordering. However, individuals feature heterogeneous pairwise relationships and interdependencies that can be adequately captured by a graph. In this setting, nodes represent subject-specific shapes, while edge weights can be used to encode similarities between subjects potentially integrating auxiliary, phenotypic information.

Semi-supervised classification Based on the graph representation, the transductive inference problem can be formulated as semi-supervised node classification, where labels are only given for nodes corresponding to subjects from the training set. As classifier we construct a multi-layer, feed-forward graph convolutional network with possibly several hidden layers each followed by a rectified linear unit (ReLU). The final layer has as many output channels as the desired number of classes and is equipped with a node-wise soft-max activation. As loss function, a cross-entropy term for each node in the training set is used. Since the model is conditioned on the adjacency of the graph, there is no need for explicit graph-based regularization: The gradient information of the loss is propagated through the model enabling it to learn representations of both labeled and unlabeled nodes. For our architecture, we opt for a spectral generalization of graph convolutions, which are based on a K^{th} -order approximation in terms of Chebyshev polynomials (Defferrard et al., 2016) and provide fast localized graph convolutions with constant learning complexity.

Node features For use as network features, the shape differences in Eqs. (1) and (2) need to be linearized. To this end, we note that the area-based differences belong to the space of symmetric positive-definite (SPD) matrices. This space again exhibits a rich, geometric structure. Remarkably, it can be equipped with a Lie group structure inducing an efficient Riemannian metric (Arsigny et al., 2006). Furthermore, the group structure provides a canonical linearization in terms of the group logarithm that we can readily employ. In particular, linear computations on the logarithms correspond to Riemannian operations on the SPD manifold and, hence, do not suffer from defects like the *swelling effect* inherent to flattening of SPD matrices. Additionally, we note that the conformal-based differences can be transformed to symmetric matrices via the left action of Λ_0 . We therefore also employ the group logarithm on the resulting operators as linearized features.

Graph construction While the construction of the FMN is driven by geometric considerations (viz. establishing consistent, group-wise correspondences), it does not encode phenotypic relationships and interdependencies between the subjects that are informative to the grading task. We therefore condition our model on another graph that also leverages




Figure 2: Illustration of two MRI-derived hippocampi from the ADNI dataset. Note the topological noise present in the right instance.

non-geometric information. Following (Parisot et al., 2018), we define the adjacency matrix W of the population graph by

$$W_{ij} = Sim(\mathcal{S}_i, \mathcal{S}_j) \sum_k \delta(m_i^k, m_j^k),$$

where Sim gauges similarity between subject shapes and δ is a threshold function testing for closeness of phenotypic measures m^k such as age and sex. Whereas the exact choice of Sim should be application-dependent, a canonical candidate is to employ a radial basis function kernel based on distances of the node features.

3. Application to Alzheimer’s Classification

Morbus Alzheimer is a complex, multifactorial, neurodegenerative disease (McGirr et al., 2020). In 2014, it was the sixth leading cause of death in the United States, accounting for 3.6% of all deaths (Taylor et al., 2017; Heron, 2016; Alzheimer’s Association et al., 2017). Moreover, nearly everyone in the final stages of Alzheimer’s disease needs constant care as the result of functional and cognitive declines (Alzheimer’s Association et al., 2017). This poses an enormous burden for the society, clearly indicating the need for a deeper understanding of the disease. Alzheimer’s disease is known to affect large regions of the human brain (Castellani et al., 2010), among others, the two hippocampi (Van Hoesen and Hyman, 1990), that play a certain role in the formation of memories (Squire, 1992).

Within the following experiment we focus on (global) morphological changes of the *right* hippocampus in order to classify the presence of Alzheimer’s disease.

3.1. Data

We employ a dataset consisting of right hippocampi (see, e.g., Figure 2) from the Alzheimer’s Disease Neuroimaging Initiative (ADNI)¹ database comprising 60 subjects showing Alzheimer’s disease and 60 cognitive normal controls. We prepared this dataset using imaging data that contains, among others, 1632 brain MRI scans, with respective hippocampus

1. adni.loni.usc.edu




Figure 3: The proposed GCN approach achieves the highest average classification accuracy of 0.787 ± 0.055 followed by CNN (0.77 ± 0.06), MCNN (0.767 ± 0.038) and MLP (0.749 ± 0.074).

segmentations, collected at four different time points. The dataset was randomly assembled from isosurfaces extracted from given connected segmentation masks at baseline time point. This experiment was first carried out in the work by [Ambellan et al. \(2021b\)](#) and a list of employed (unique) scan ids can be found there. Public availability of the utilized hippocampus segmentations as part of the ADNI database enables reproducibility of the experiment below.

3.2. Disease grading

We evaluate the proposed Graph Convolutional Network (GCN) for the discrimination between normal controls and subjects with Alzheimer’s disease. To this end, we employ three layers of second-order graph convolutions with input dimensions $(n_{\text{shapes}}, 64, 64)$. As phenotype measures, we selected sex and Apolipoprotein E (APOE) genotype as proposed in ([Parisot et al., 2018](#)). Gene APOE appears in three major types: E2, E3 and E4. Especially E4 is known as genetic risk factor for Alzheimer’s disease. Affinity between two connected nodes is considered high if sex and APOE type coincide. Regarding classification accuracy we compare it to three different approaches. On the one hand, we employ classification approaches using a MultiLayer Perceptron (MLP), as well as a Convolutional Neural Network (CNN) applied on the area- and conformal-based shape differences with isophotic metric, as detailed in ([Huang et al., 2019](#)). On the other hand, we report results for MeshCNN (MCNN) as described by [han](#). In contrast to all other methods that were taken into consideration, MCNN is applied directly on the shape-forming triangular mesh. Note that MLP, CNN and MCNN are following the *inductive* learning approach. We eval-

uated all methods on a 70%/30% training/testing split performing a stratified Monte Carlo cross-validation drawing 300 times for MLP, CNN, GCN and due to its demanding complexity 10 times for MCNN. The results are summarized in Figure 3, indicating that (i) the transductive GCN approach achieves the highest (78.7%) average classification accuracy, (ii) CNN and MCNN are approximately on par (77.0%/76.7%), and (iii) MLP achieves the lowest performance with 74.9% accuracy.

4. Conclusion and future work

In this work, we presented a geometric classification scheme from shape data employing a flexible, yet descriptive characterization of shape variability. Based on a graph convolutional neural network, we further perform transductive inference taking the irregular structure in sampling patterns of clinical data sets into account. Furthermore, we extended the functional characterization of shape variation via an alternative metric that is sensitive to extrinsic curvature and employed a geometric linearization based on the Log-Euclidean framework for positive matrices.

In application to Alzheimer’s disease classification, we achieved an improved classification that outperforms recent work for deep learning on 3D surfaces (han) as well as inductive inference from functional descriptions (Huang et al., 2019). Furthermore, our method is able to significantly decrease the gap towards Riemannian shape spaces that rely on the more restrictive setting of dense vertex correspondence (Ambellan et al., 2021b) and for which 80.4% classification accuracy has been reported. In future work, we plan to investigate further the potential of our approach for topologically-varying and incomplete shape collections in order to widen to scope of shape analysis methodology and to provide more extensive empirical evidence of its performance.

As the right hippocampus shape very likely does not provide a complete enough picture to classify Alzheimer’s disease, we probably have already reached a certain limit in our experimental setup. Therefore, on the application side, a promising line of future work is to extend the dataset with the respective left hippocampi. This will help to assess the very complex Alzheimer’s disease in a more anatomically holistic manner. Another direction for future work is the graph construction itself, which should be further analyzed as it poses a structural core element in our setup and potentially is the key property for discriminative tasks such as disease grading.

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689) and by the Bundesministerium für Bildung und Forschung (BMBF) through The Berlin Institute for the Foundations of Learning and Data (BIFOLD) - (ref. 01IS18025A and ref. 01IS18037A). Furthermore, this work relies on data of the Alzheimer’s Disease Neuroimaging Initiative (ADNI)².

2. Data collection and sharing for this project was funded by the ADNI (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and

References

meshcnn.

Alzheimer’s Association et al. 2017 Alzheimer’s disease facts and figures. *Alzheimer’s & Dementia*, 13(4):325–373, 2017.

Felix Ambellan, Hans Lamecker, Christoph von Tycowicz, and Stefan Zachow. Statistical shape models: understanding and mastering variation in anatomy. In *Biomedical Visualisation*, pages 67–84. Springer, 2019.

Felix Ambellan, Stefan Zachow, and Christoph von Tycowicz. Geodesic b-score for improved assessment of knee osteoarthritis. In *International Conference on Information Processing in Medical Imaging*, pages 177–188. Springer, 2021a.

Felix Ambellan, Stefan Zachow, and Christoph von Tycowicz. Rigid motion invariant statistical shape modeling based on discrete fundamental forms: Data from the Osteoarthritis Initiative and the Alzheimer’s disease Neuroimaging Initiative. *Medical Image Analysis*, 73:102178, 2021b.

Pierre-Louis Antonsanti, Joan Glaunès, Thomas Benseghir, Vincent Jugnon, and Irène Kaltenmark. Partial matching in the space of varifolds. In *International Conference on Information Processing in Medical Imaging*, pages 123–135. Springer, 2021.

Vincent Arsigny, Pierre Fillard, Xavier Pennec, and Nicholas Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magn Reson Med*, 56(2):411–421, 2006.

Martin Bauer, Martins Bruveris, and Peter W Michor. Overview of the geometries of shape spaces and diffeomorphism groups. *Journal of Mathematical Imaging and Vision*, 50(1): 60–97, 2014.

Mario Botsch, Leif Kobbelt, Mark Pauly, Pierre Alliez, and Bruno Lévy. *Polygon Mesh Processing*. Natick, Mass.: AK Peters, Ltd., 2010.

Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond Euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer’s Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

- Rudy J Castellani, Raj K Rolston, and Mark A Smith. Alzheimer disease. *Disease-a-month: DM*, 56(9):484, 2010.
- Etienne Corman, Justin Solomon, Mirela Ben-Chen, Leonidas Guibas, and Maks Ovsjanikov. Functional characterization of intrinsic and extrinsic geometry. *ACM Transactions on Graphics (TOG)*, 36(2):1–17, 2017.
- M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pages 3844–3852, 2016.
- Manfredo P. do Carmo. *Differential geometry of curves and surfaces*. Prentice Hall, 1976.
- Guido Gerig, James Fishbaugh, and Neda Sadeghi. Longitudinal modeling of appearance and shape and its potential for clinical use. *Medical Image Analysis*, 33:114–121, 2016.
- Melonie Heron. Deaths: Leading causes for 2014. *Natl Vital Stat Rep*, 65(5):1–96, 2016.
- Ruqi Huang, Panos Achlioptas, Leonidas Guibas, and Maks Ovsjanikov. Limit shapes—a tool for understanding shape differences and variability in 3d model collections. In *Computer Graphics Forum*, volume 38, pages 187–202. Wiley Online Library, 2019.
- Samantha McGirr, Courtney Venegas, and Arun Swaminathan. Alzheimers disease: A brief review. *Journal of Experimental Neurology*, 1(3):89–98, 2020.
- Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019.
- Esfandiar Nava-Yazdani, Hans-Christian Hege, and Christoph von Tycowicz. A hierarchical geodesic model for longitudinal analysis on manifolds. *Journal of Mathematical Imaging and Vision*, 64(4):395–407, 2022.
- Rune Kok Nielsen, Sune Darkner, and Aasa Feragen. TopAwaRe: Topology-aware registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 364–372. Springer, 2019.
- Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)*, 31(4):1–11, 2012.
- Sarah Parisot, Sofia Ira Ktena, Enzo Ferrante, Matthew Lee, Ricardo Guerrero, Ben Glocker, and Daniel Rueckert. Disease prediction using graph convolutional networks: Application to Autism Spectrum Disorder and Alzheimer’s disease. *Medical Image Analysis*, 48:117–130, 2018.
- Helmut Pottmann, Tibor Steiner, Michael Hofer, Christoph Haider, and Allan Hanbury. The isophotic metric and its application to feature sensitive morphology on surfaces. In *European Conference on Computer Vision*, pages 560–572. Springer, 2004.

- Martin Reuter, Franz-Erich Wolter, and Niklas Peinecke. Laplace-Beltrami spectra as ‘Shape-DNA’ of surfaces and solids. *Computer Aided Design*, 38(4):342–366, 2006.
- Raif M. C. Rustamov, M. Ovsjanikov, Omri Azencot, Mirela Ben-Chen, and Frédéric Chazal. Map-based exploration of intrinsic shape differences and variability. *ACM Trans. Graphics*, 32(4)(72):1–12, 2013.
- Larry R Squire. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychological review*, 99(2):195, 1992.
- Alexander Tack, Felix Ambellan, and Stefan Zachow. Towards novel osteoarthritis biomarkers: Multi-criteria evaluation of 46,996 segmented knee MRI data from the Osteoarthritis Initiative. *PloS one*, 16(10):e0258855, 2021.
- Christopher A Taylor, Sujay F Greenlund, Lisa C McGuire, Hua Lu, and Janet B Croft. Deaths from Alzheimer’s Disease—United States, 1999–2014. *Morbidity and Mortality Weekly Report*, 66(20):521, 2017.
- Gary W Van Hoesen and Bradley T Hyman. Hippocampal formation: anatomy and the patterns of pathology in Alzheimer’s disease. *Progress in brain research*, 83:445–457, 1990.
- Christoph von Tycowicz. Towards shape-based knee osteoarthritis classification using graph convolutional networks. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 750–753. IEEE, 2020.
- Christoph von Tycowicz, Felix Ambellan, A Mukhopadadhyay, and Stefan Zachow. An efficient Riemannian statistical shape model using differential coordinates: With application to the classification of data from the Osteoarthritis Initiative. *Med. Image. Anal.*, 43:1–9, 2018.
- Fan Wang, Qixing Huang, and Leonidas J Guibas. Image co-segmentation via consistent functional maps. In *Proceedings of the IEEE international conference on computer vision*, pages 849–856, 2013.