3D Shape Analysis: From Classical Optimisation Methods to Feature Learning for Shape Matching

N. El Amrani[1] L. Bastian[2,3] V. Ehm[2,3] Z. Lähner[1,4] F. Bernard[1]

[1]University of Bonn [2] Technical University of Munich [3]Munich Center for Machine Learning [4]Lamarr Institute

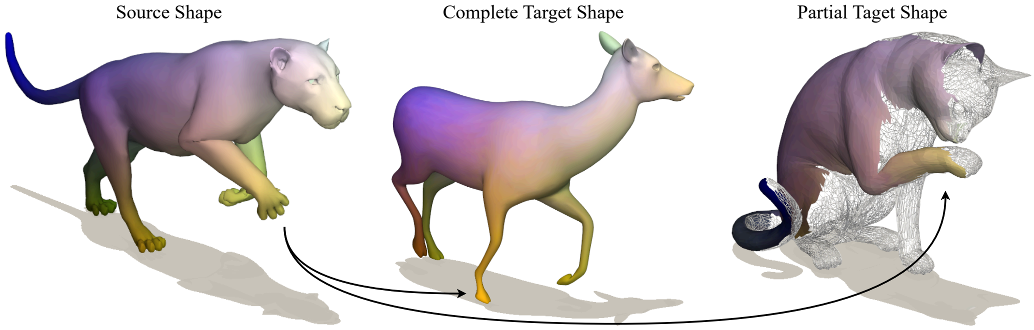

Figure 1: **Illustration of 3D shape matching:** Given a source shape, the task is to find for each point on the source shape the corresponding on the target shape. For partial shape matching, the task has the additional objective of predicting which part is present in the source shape.

Abstract

The field of 3D shape analysis is concerned with the extraction of "useful" information from geometric data. Shape analysis has a high relevance for a wide range of applications, such as autonomous driving, biomedicine, or augmented/virtual reality. A core task of 3D shape analysis is shape matching, i.e. identifying correspondences between given shapes. While traditional shape matching methods rely on optimising a task-specific objective function, modern shape matching oftentimes involves data-driven components. We will first introduce traditional methods for shape matching, starting with the linear assignment problem and the quadratic assignment problem. We then present product graph formalisms in different settings, including 2D to 2D, 2D to 3D or shape to image, and 3D to 3D shape matching. We then discuss recent developments in learning-based shape correspondence methods, from learning shape correspondence with topological data structures to spectral approaches that provide efficient structure and circumvent annotations altogether. Furthermore, we discuss the practical relevance of these methods to application domains in image-to-image and shape-to-image correspondence, medical imaging and surgical navigation, and discuss how recent developments in foundation models play a role in shape analysis. Finally, the tutorial will conclude by addressing the challenges of shape matching, including handling partial shapes, and will explore potential future directions in the field.

1. Prerequisites & Intended Audience

This tutorial is designed for various backgrounds from graduate students to researchers from both academia and industry. While the primary focus of the tutorial is on shape matching within the broader field of shape analysis, the insights and knowledge shared are valuable to participants from various disciplines. We assume a general knowledge about 3D shapes; however, no prior experience with shape analysis or matching is required, as all relevant concepts will be thoroughly explained during the tutorial.

2. Outline tutorial

- **20 min: Introduction, motivation and overview** (*Florian Bernard*) Shape analysis has a wide range of applications, including sketch-based shape retrieval [ERB*12], shape segmentation and clustering [Läh21], feature computation [GRE*23], statistical shape analysis [EACB24], and understanding deformations and shape interpolation [CEEA*24, ENK*21, EC20]. A fundamental approach to achieve these applications is by solving the **shape matching** problem, which is the primary focus of this tutorial. In shape matching, the goal is to identify semantically corresponding parts between shapes. Solution strategies include various paradigms, e.g. spectral methods and assignment problems, among others. The tutorial will provide an overview of these key paradigms and their relevance to shape analysis.
- **60 min: Optimisation-based shape matching** (*Florian Bernard and Viktoria Ehm*) We focus on axiomatic methods that build on various forms of assignment problems. We begin by exploring classical assignment problems, such as the linear and the quadratic assignment problem. Further, we introduce product graphs, which are well-known from the classical Dynamic Time Warping problem [SC78]. These product graphs are also used for 2D-to-2D shape matching and 2D-to-3D shape matching [LRS*16]. In contrast to these cases, where matching can be represented as a one-dimensional path, 3D shape matching requires a more complex formulation. Specifically, it is modeled as a constrained integer linear program [WSSC11]. We further explain how the algorithm can be extended to the challenging problem of partial shape matching [ERE*23, EGR*24]. Finally, we show how 2D-to-3D product graph can be utilised for 3D-to-3D shape matching [RB24].
- **10 min: Q&A**
- **20 min: Break**
- **30 min: Deep shape matching and applications** (*Lennart Bastian*) Shape matching has distinguished itself from other correspondence problems in computer vision because shapes can be treated as manifold – a strong prior which has given rise to powerful frameworks enabling unsupervised matching. In this part, we will dive into what distinguishes this domain. Manifold assumptions can be leveraged to construct descriptors [SOG09, ASC11] or as priors for learning [SACO22]. Furthermore, they have given rise to a powerful framework of functional maps [OBCS*12], which reduce the challenging matching problem to a linear change in a spectral basis, thanks to a simplifying manifold assumption [MRR*19]. In a learned setting, this manner of spectral matching can be differentiated through, which provides a stable supervision signal when provided appropriate structure [LRR*17, DSO20, RPWO19]. This has given rise to many applications, including, for example, statistical shape modeling in medical imaging [EACB24, BBH*23], or modeling anatomical morphology [MBT*23, LVdF*24], where data-driven descriptors are essential to handle noisy or incomplete surface representations.
- **5 min: Q&A**
- **50 min: Challenges and future directions** (*Zorah Lähner and Nafie El Amrani*)
  In this section of the tutorial, we will discuss the key challenges that remain unsolved in the field of 3D shape learning, which can be broadly categorized into two main areas: scalability and severe non-isometry of shapes. Scalability challenges arise from several factors. One issue is the high or significantly varying resolution of shapes, which complicates efficient processing. Additionally, instead of simply matching shapes in pairs, achieving cycle consistency when matching multiple shapes simultaneously presents a more complex problem. Another important aspect is the ability of shape matching methods to generalize across different object types, such as natural and man-made objects, which often have distinct characteristics. The second challenge, severe non-isometry, refers to the difficulties in matching shapes that are only partially available, as shown in Fig 2, or in matching shapes that have topological differences. Looking ahead to future directions, we propose a first step towards addressing these scalability issues by introducing larger and more scalable datasets. In particular, we present the BeCoS benchmark [EEAX*24], which enables the generation of new shapes, helping to push the boundaries of 3D shape learning.
- **5 min: Q&A**

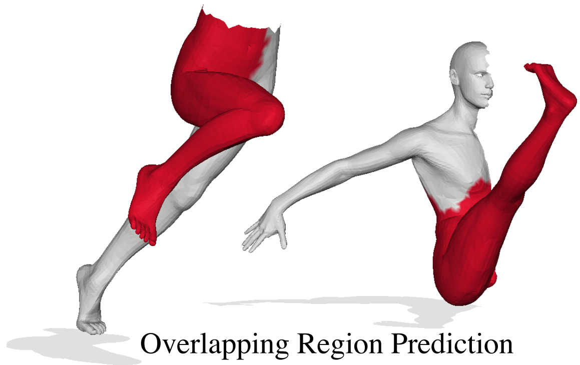

Figure 2: **Partial shape matching** has the additional challenge of (explicitly or implicitly) identifying the **overlapping region** between the two input partial shapes. Figure taken from [EGR*24].

3. Tutorial's format

- 2×90 slots (half a day)
- Lecture style

4. Tutorial's History

- Venue: GCPR/VMV 2024 in Munich
- Date: Tuesday 10th of September 2024
- Number of attendees: ≈75

The first iteration of the tutorial successfully raised awareness of the growing field of shape analysis and matching, with around 75 attendees at the GCPR-VMV conference in Munich and positive feedback. In the coming iterations, we hope to share expertise gathered over the years with the broader international communities, specifically in computer graphics and vision, to help foster continued methodological development and translation into various application domains.

5. Presenters

Nafie El Amrani.

- Institution: University of Bonn
- URL: https://nafieamrani.github.io/

Nafie is a PhD student at the University of Bonn under the supervision of Prof. Florian Bernard in the Learning and Optimisation for Visual Computing group. Previously, he received his Master's degree in Computer Science (in 2023) as well as his Bachelor's degree in Computer Science & Engineering (in 2021) at Delft University of Technology (TU Delft). His research interests are multimodal shape analysis, computer vision and deep learning.

Lennart Bastian.

- Institution: Technical University of Munich
- URL: https://www.cs.cit.tum.de/camp/members/lennart-bastian/

Lennart Bastian is a PhD student at the Technical University of Munich (TUM), where he works under the supervision of Prof. Dr. Nassir Navab and coordinates the Surgical Data Science team. His research focuses on the intersection of 3D geometry in machine learning and its applications in medical imaging and computer-assisted interventions, with particular emphasis on the analysis of motion, deformation, and shape. Lennart completed his M.Sc. in Applied Mathematics at TUM, where he wrote his thesis on storage efficient semi-definite programs under Prof. Dr. Michael Ulbrich in the optimization department. He earned his Bachelor's degree in Mathematics and Computer Science from the New York University Courant Institute.

Viktoria Ehm.

- Institution: Technical University of Munich
- URL: https://vikiehm.github.io

Viktoria is a PhD student in the Computer Vision Group at the TU Munich, supervised by Prof. Daniel Cremers. She received her bachelor's degree at Ludwig Maximilans University (LMU) and her master's degree at Technical University of Munich (TUM), both in computer science. Her research interests mainly lie in the area of 3D shape matching, while she is currently focusing on partial shape matching in both the classical optimisation and the deep learning domain.

Zorah Lähner.

- Institution: University of Bonn & Lamarr Institute
- URL: https://geometryinml.cs.uni-bonn.de

Zorah is an assistant professor (tenure track) and head of the "Geometry in Machine Learning" group at the University of Bonn and the Lamarr Institute for Machine Learning and Artificial Intelligence. She received her PhD from the Technical University of Munich and was a postdoctoral researcher at the University of Siegen. She is interested in how geometric properties can be used to form effective priors and guide optimisation processes in geometric deep learning and 3D computer vision applications. For example, choosing a different representation can make a huge difference in the capabilities and efficiency of geometric data, particularly in non-Euclidean domains.

Florian Bernard.

- Institution: University of Bonn
- URL: https://lovc.cs.uni-bonn.de/

Florian Bernard is Associate Professor at the University of Bonn, where he heads the Learning and Optimisation for Visual Computing group. Before joining University of Bonn, he was Visiting Professor at the chair of Computer Vision & Artificial Intelligence at the Technical University of Munich. He also held a position as postdoctoral researcher in the Graphics, Vision and Video group at the Max-Planck-Institute for Informatics, as well as at the University of Luxembourg, where he also received his Ph.D. degree. His research lies at the intersection of visual computing, mathematical optimisation and machine learning, and aims to understand commonalities in large collections of visual data.

References

[ASC11] Aubry M., Schlickewei U., Cremers D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)* (2011), IEEE, pp. 1626–1633. 2

[BBH*23] Bastian L., Baumann A., Hoppe E., Bürgin V., Kim H. Y., Saleh M., Busam B., Navab N.: S3m: scalable statistical shape modeling through unsupervised correspondences. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2023), Springer, pp. 459–469. 2

[CEEA*24] Cao D., Eisenberger M., El Amrani N., Cremers D., Bernard F.: Spectral meets spatial: Harmonising 3d shape matching and interpolation. In *CVPR* (2024). 2

[DSO20] Donati N., Sharma A., Ovsjanikov M.: Deep geometric functional maps: Robust feature learning for shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 8592–8601. 2

[EACB24] El Amrani N., Cao D., Bernard F.: A universal and flexible framework for unsupervised statistical shape model learning. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)* (Oct. 2024). 2

[EC20] Eisenberger M., Cremers D.: Hamiltonian dynamics for real-world shape interpolation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16* (2020), Springer, pp. 179–196. 2

[EEAX*24] Ehm V., El Amrani N., Xie Y., Bastian L., Gao M., Wang W., Sang L., Cao D., Lähner Z., Cremers D., Bernard F.: Beyond complete shapes: A quantitative evaluation of 3d shape matching algorithms, 2024. URL: https://arxiv.org/abs/2411.03511, arXiv:2411.03511. 2

[EGR*24] Ehm V., Gao M., Roetzer P., Eisenberger M., Cremers D., Bernard F.: Partial-to-partial shape matching with geometric consistency. *CVPR* (2024). 2

[ENK*21] Eisenberger M., Novotny D., Kerchenbaum G., Labatut P., Neverova N., Cremers D., Vedaldi A.: Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 7473–7483. 2

[ERB*12] Eitz M., Richter R., Boubekeur T., Hildebrand K., Alexa M.: Sketch-based shape retrieval. *ACM Transactions on graphics (TOG) 31*, 4 (2012), 1–10. 2

[ERE*23] Ehm V., Roetzer P., Eisenberger M., Gao M., Bernard F., Cremers D.: Geometrically consistent partial shape matching. *International Conference on 3D Vision (3DV)* (2023). 2

[GRE*23] Gao M., Roetzer P., Eisenberger M., Lähner Z., Moeller M., Cremers D., Bernard F.: SIGMA: Quantum scale-invariant global sparse shape matching. In *International Conference on Computer Vision (ICCV)* (2023). 2

[Läh21] Lähner Z.: *Continuous Correspondence of Non-Rigid 3D Shapes*. PhD thesis, Technische Universität München, 2021. 2

[LRR*17] Litany O., Remez T., Rodola E., Bronstein A., Bronstein M.: Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE international conference on computer vision* (2017), pp. 5659–5667. 2

[LRS*16] Lahner Z., Rodola E., Schmidt F. R., Bronstein M. M., Cremers D.: Efficient globally optimal 2d-to-3d deformable shape matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016), pp. 2185–2193. 2

[LVdF*24] Leblanc L., Vialle R., de Farias C., Saghbini E., Marturi N., Tamadazte B.: Pedicle drilling planning transfer for spine surgery using functional map correspondences. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2024)* (2024). 2

[MBT*23] Magnet R., Bloch K., Taverne M., Melzi S., Geoffroy M., Khonsari R. H., Ovsjanikov M.: Assessing craniofacial growth and form without landmarks: A new automatic approach based on spectral methods. *Journal of Morphology 284*, 8 (2023), e21609. 2

[MRR*19] Melzi S., Ren J., Rodolà E., Sharma A., Wonka P., Ovsjanikov M.: Zoomout: spectral upsampling for efficient shape correspondence. *ACM Trans. Graph. 38*, 6 (Nov. 2019). URL: https://doi.org/10.1145/3355089.3356524, doi:10.1145/3355089.3356524. 2

[OBCS*12] Ovsjanikov M., Ben-Chen M., Solomon J., Butscher A., Guibas L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG) 31*, 4 (2012), 30:1–30:11. 2

[RB24] Roetzer P., Bernard F.: Spidermatch: 3d shape matching with global optimality and geometric consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 14543–14553. 2

[RPWO19] Ren J., Panine M., Wonka P., Ovsjanikov M.: Structured regularization of functional map computations. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 39–53. 2

[SACO22] Sharp N., Attaiki S., Crane K., Ovsjanikov M.: Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (TOG) 41*, 3 (2022), 1–16. 2

[SC78] Sakoe H., Chiba S.: Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing* (1978). 2

[SOG09] Sun J., Ovsjanikov M., Guibas L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum* (2009), vol. 28, Wiley Online Library, pp. 1383–1392. 2

[WSSC11] Windheuser T., Schlickewei U., Schmidt F. R., Cremers D.: Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *2011 International Conference on Computer Vision* (2011), IEEE, pp. 2134–2141. 2