Statement of Purpose

Aditya Chetan (Computer Science Ph.D. applicant for Fall 2022)

I want to pursue a Ph.D. in Computer Science at the University of Washington to conduct research in **data-driven pipelines for Geometry Processing**. Data-centric frameworks have driven significant progress in image-processing as they have proved highly effective in modeling semantic patterns and encoding task compositionality. In recent years, these methods have also been increasingly adopted for 3D data, replacing traditional algorithms for Geometry Processing. I have immersed myself in exploring and understanding such approaches that adopt this data-centric mindset for Geometry Processing.

Research Background. I completed my undergraduate studies finishing at the **top of my cohort** from **IIIT Delhi** with a major in **Computer Science and Applied Mathematics**. Initially, I became interested in developing novel data-driven systems for social network analysis and natural language processing (NLP). My thesis explored a special case of anomaly detection in Online Social Networks with Prof. Tanmoy Chakraborty. This work was accepted at top-tier venues in the areas of Social Network Analysis and Data Mining, like **ASONAM'18** [5], **ACM TIST** [1], and **WSDM'19** [4], where I was also a **primary author**. We also received the **Innovative Student Projects Award** from the Indian National Academy of Engineering for this work, **one of the highest recognitions** for undergraduates in India. I led every aspect of this work: ideation, writing code, empirical analysis, to writing compelling manuscripts and rebuttals; skills which I feel have armed me to succeed as a graduate student. I also took a semester off from my studies to visit the **Max Planck Institute for Informatics**, where I worked on information extraction from unstructured text with Prof. Andrew Yates and Dr. Paramita Mirza. Currently, I am working at **Microsoft Research India** (MSR) as part of the **highly competitive Research Fellows** program, which aims to prepare research-oriented undergraduates for higher studies.

In recent years, my interests have evolved towards data-driven methodologies for Geometry Processing. I gained the technical expertise required for research in this area by undertaking courses and projects in **Scientific Computing**, **Numerical PDEs**, and **Deep Learning**. In my junior year, I undertook an independent study in **Differential Geometry** with Prof. Kaushik Kalyanaraman that culminated in the implementation of algorithms for computing discrete curvature [2]. At the **Fields Institute, Toronto**, I worked with Prof. Alec Jacobson at the University of Toronto on **Geometry Processing for Signed Distance Fields** through the **Fields Undergraduate Summer Research Program (FUSRP) 2021**. At **MSR**, I am working on improving the **adversarial robustness of neural networks** using ideas from **topology**. Below I will summarise specific experiences in Geometry Processing and learning-based methods and outline how I believe data-centric approaches can be integrated into Geometry Processing research.

Representations for Geometry. While working with Prof. Jacobson's group, I recognized that one of the primary challenges for adopting data-driven approaches for Geometry Processing is selecting a representation for 3D data. Multiple choices exist, such as polygonal meshes, point clouds, parametric functions, and so on. Recently, (Signed) Distance Fields (SDFs) have emerged as a popular representation that fits well into learning-based frameworks [9, 10] and offers advantages like ease of iterative geometric modeling and shape editing. However, many standard Geometry Processing operations (such as FEM-based analysis) under-utilize SDFs, as they require extracting an explicit mesh from them, which increases performance costs. To address this challenge in part, we developed algorithms to perform Geometry Processing directly on SDFs. Specifically, we focussed on Mean Curvature Flow (MCF) and proved that applying Gaussian Smoothing on the SDF results in MCF on the surface under certain conditions [3]. More generally, we found that to develop algorithms that operate directly on SDFs, preserving the characteristic properties of the SDF (like continuity and unit-norm gradient) at each step is imperative and not achieved by existing works. Currently, I am focusing on a learning-based system that facilitates this by converting approximate-SDFs (that may not satisfy all properties of SDFs) to exact SDFs. I believe this approach will enable the development of algorithms for performing Geometry Processing operations beyond MCF directly on SDFs. More broadly, in my Ph.D., I would like to work on developing algorithms that adopt a data-driven approach to Geometry Processing for various surface representations.

Adversarial Geometry. Another significant challenge for data-driven systems is the ability to handle noisy geometric data that is common in the wild. For this, I believe that we need to adopt robust learning algorithms that can handle outliers, adversarial examples, and fine-grained surface details. At MSR with Dr. Nipun Kwatra, I am developing such adversarially robust learning algorithms for training neural network classifiers. We factor in the geometry of the underlying data manifold by training models to learn a smooth function of the distance of an input point from the manifold. We reason that doing so provides a richer signal for the model to learn rather than learning class labels, and this improves robustness to adversarial samples. We can also label points predicted to lie beyond a threshold as outliers. As the lead investigator of this project, I have been responsible for its conceptualization along with extensive empirical experimentation. We have achieved promising results on synthetic datasets and are now scaling our method to real-world datasets like MNIST and CIFAR-10. Though we proposed our method primarily for images on a 2D pixel grid, I believe that we can adapt it to learn robust Neural SDFs [10] for 3D surfaces too. I believe this approach will benefit applications like mesh reconstruction from point clouds, where distinguishing between on-manifold samples and outliers is critical. During my Ph.D., I would also like to focus on developing techniques for robust data-driven Geometry Processing.

***Traditional vs. Neural methods*.** Lastly, in my view, to truly benefit from learning-based algorithms, we need to thoroughly explore the limits of hyper-optimized traditional algorithms and use them as benchmarks to design new data-driven methods. With Prof. Jacobson and his students Silvia Sellán and Sarah Kushner, I am working on a project in this spirit on computing volumes swept by surfaces in motion (Swept Volumes). Though Silvia's recent work [11] does this effectively using SDFs, it cannot handle cases where the surface undergoes non-rigid deformations during motion. To solve this issue, I am working on an optimized version of the traditional Bounding Volume Hierarchy (BVH) data structure that would store the SDF in 4D space-time [8] instead of 3D space. Our approach has controllable sources of error, like the resolution of the grid storing the SDF. In contrast, learning a Neural Implicit [10] for the Swept Volume gives us a continuous representation independent of grid resolution but less controllable error. As a part of my Ph.D., I would like to find the right balance between traditional and neural methods for integrating them in Geometry Processing pipelines.

***Teaching & Outreach*.** My decision to pursue a Ph.D. is guided by my inclination towards a career in academia, as it promises a unique blend of teaching and mentorship opportunities along with research. I had a positive teaching experience as the only **teaching assistant** of a **graduate-level course** at IIIT Delhi. Since this was the course's first offering, not only did I mentor students in projects and grade assignments but also helped with content development. I received **positive feedback** for this role which gave me confidence in my career goals. To attract and nurture new talent in our community, I am also a part of the **SIGGRAPH Research Career Development Committee** led by Prof. Justin Solomon (MIT). I have been focusing on curating resources to introduce new students to graphics research. I am also committed to open science through high-quality, open-source software. Open-source implementations of my projects have been **well received** by the community (aggregating **250+ stars** on GitHub). I plan to continue releasing open-source toolkits for data-driven Geometry Processing during my Ph.D. for the benefit of the broader graphics and artist community.

Why UW? At UW, I am interested in the work of **Prof. Adriana Schulz** on Computational Design. I am excited about the recent work of Prof. Schulz's group on using domain-specific languages that enable benefits like bi-directional editing between figures and programs or joint optimization of design and fabrication constraints, among others. I am also interested in exploring neuro-symbolic approaches for computational design [6, 7] that satisfy some required geometric or application-based constraints [12]. Given my interests in data-driven methods for Geometry Processing and my past experiences in Natural Language Processing, I would like to explore ways to exploit recent advances in learning-based program synthesis for domain-specific languages in the context of computational design. I also look forward to collaborating with other members of the UW Graphics and Imaging Laboratory (GRAIL) and the Digital Fabrication Center (DFab) on problems of mutual interest.

I am confident that I would fit well in the vibrant student community at UW and am looking forward to engaging in productive discussions with my peers from diverse backgrounds. I believe that I bring a unique background to UW, and I am confident that pursuing a Ph.D. there will provide me with the experience and expertise to pursue my long-term career goals of a research career in academia.

References

[1] Arora, U., Dutta, H. S., Joshi, B., Chetan, A., and Chakraborty, T. Analyzing and detecting collusive users involved in blackmarket retweeting activities. *ACM Trans. Intell. Syst. Technol. 11*, 3 (Apr. 2020).

[2] Chetan, A. DiffGeoOps: Python library implementing algorithms for computing discrete curvatures. GitHub repository available at https://github.com/justachetan/DiffGeoOps (Accessed November 20 2021).

[3] Chetan, A., Guo, J., Mattos Da Silva, L., Ridgway, J., Sellán, S., and Jacobson, A. The Space of Signed Distance Fields. Available at http://www.fields.utoronto.ca/sites/default/files/FUSRP_2021__Project_3__Final%20Report.pdf (Accessed November 24 2021). Fields Undergraduate Summer Research Program 2021 Project Report.

[4] Chetan, A., Joshi, B., Dutta, H. S., and Chakraborty, T. Corerank: Ranking to detect users involved in blackmarket-based collusive retweeting activities. In *WSDM 2019*.

[5] Dutta, H. S., Chetan, A., Joshi, B., and Chakraborty, T. Retweet us, we will retweet you: Spotting collusive retweeters involved in blackmarket services. In *ASONAM 2018*.

[6] Jones, R. K., Barton, T., Xu, X., Wang, K., Jiang, E., Guerrero, P., Mitra, N. J., and Ritchie, D. Shapeassembly: Learning to generate programs for 3d shape structure synthesis. *ACM Transactions on Graphics (TOG), Siggraph Asia 2020 39*, 6 (2020), Article 234.

[7] Jones, R. K., Charatan, D., Guerrero, P., Mitra, N. J., and Ritchie, D. Shapemod: Macro operation discovery for 3d shape programs. *ACM Transactions on Graphics (TOG), Siggraph 2021 40*, 4 (2021), Article 153.

[8] Klosowski, J. T., Held, M., Mitchell, J. S. B., Sowizral, H., and Zikan, K. Efficient collision detection using bounding volume hierarchies of k-dops. *IEEE Transactions on Visualization and Computer Graphics 4*, 1 (jan 1998), 21–36.

[9] Mescheder, L. M., Oechsle, M., Niemeyer, M., Nowozin, S., and Geiger, A. Occupancy networks: Learning 3d reconstruction in function space. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019), 4455–4465.

[10] Park, J. J., Florence, P. R., Straub, J., Newcombe, R. A., and Lovegrove, S. DeepSDF: Learning continuous signed distance functions for shape representation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019).

[11] Sellán, S., Aigerman, N., and Jacobson, A. Swept volumes via spacetime numerical continuation. *ACM Transactions on Graphics* (2021).

[12] Zhao, A., Xu, J., Konaković-Luković, M., Hughes, J., Spielberg, A., Rus, D., and Matusik, W. Robogrammar: Graph grammar for terrain-optimized robot design. *ACM Trans. Graph. 39*, 6 (nov 2020).