Sai Praveen Bangaru

sbangaru@nvidia.com | LinkedIn: saipraveen-bangaru-2ba085a0 | people.csail.mit.edu/sbangaru

Education

Massachusetts Institute of Technology | PhD, Electrical Engineering and Computer Science
- Duration: Aug 2019 – Apr 2024
- GPA: 5.0   (5.0 scale)

Carnegie Mellon University | Master's Degree (MSCS), Computer Science Department
- Duration: Aug 2017 – Dec 2018
- QPA: 4.05   (4.3 scale)

Indian Institute of Technology Madras | Bachelor's Degree (BTech), Computer Science Department
- Duration: Aug 2013 - Jul 2017
- CGPA: 9.34   (10.0 scale)

Publications

Distributions for Compositionally Differentiating Parametric Discontinuities | OOPSLA 2024 (ACM SIGPLAN)

J. Michel, K. Mu, X. Yang, S. P. Bangaru, E. R. Collins, G. Bernstein, J. Ragan-Kelley, M. Carbin, TM. Li

Link: https://dspace.mit.edu/handle/1721.1/154393

- To be presented as a talk at *OOPSLA 2024* held in Pasadena, CA, USA, as part of the ACM SIGPLAN conference.
- This paper proposes a new compiler, Potto, to differentiate programs that contain integrals over discontinuous functions, a common pattern across computer graphics, simulation and statistics. Potto uses the theory of distributions to elegantly prove that its semantics are sound, which notably allows for the composition of sub-programs with discontinuities, paving the way for differentiating large-scale discontinuous programs.

Importance Sampling BRDF Derivatives | SIGGRAPH 2024 (ACM ToG.)

Y. Belhe, B. Xu, S. P. Bangaru, R. Ramamoorthi, TM. Li

Link: https://yashbelhe.github.io/dbrdf/

- To be presented as a talk at *SIGGRAPH 2024* held in Denver, Colorado. Submitted to ACM ToG. as a journal paper.
- This paper proposes a set of techniques to efficiently sample the derivatives of a wide range of BRDF models. The method shows how to exploit the common structure of various BRDFs to derive importance sampling distributions by using a mix of positivization & mixture decomposition techniques.

SLANG.D: Fast, Modular & Differentiable Shader Programming | SIGGRAPH Asia 2023 (ACM ToG.)

S. P. Bangaru, L. Wu, TM. Li, J. Munkberg, G. Bernstein, J. Ragan-Kelley, F. Durand, A. Lefohn, Y. He

Link: https://people.csail.mit.edu/sbangaru/projects/slang-2023/

- To be presented as a talk at *SIGGRAPH Asia 2023* held in Sydney, Australia.
- This paper presents a HLSL-like shader programming language with first-class automatic differentiation primitives that interoperate with modern language features such as generics & interfaces to enable the cross-pollination of large-scale rendering frameworks with machine learning frameworks.

Warped-Area Reparameterization of Differential Path Integrals | SIGGRAPH Asia 2023 (ACM ToG.)

Peiyu Xu, S. P. Bangaru, TM. Li, S. Zhao   *Link: https://shuangz.com/projects/psdr-was-sa23/*

- To be presented as a talk at *SIGGRAPH Asia 2023* held in Sydney, Australia.
- This paper shows how to combine the warped-area reparameterization with the material-form parameterization to present an efficient way to estimate derivatives w.r.t geometry in the presence of complex global illumination.
- Received a *Best Paper* Award.

Differentiable Rendering of Neural SDFs through Reparameterization | SIGGRAPH Asia 2022 (Conf.)

S. P. Bangaru, M. Gharbi, TM. Li, F. Luan, K. Sunkavalli, M. Hasan, S. Bi, Z. Xu, G. Bernstein, F. Durand

Link: https://people.csail.mit.edu/sbangaru/projects/dsdf-2022/

- Presented as a pre-recorded talk at *SIGGRAPH Asia 2022* held in Daegu, Republic of Korea.

- This paper shows how to differentiate a neural SDF renderer without approximations using reparameterization, with applications to 3D reconstruction on real and synthetic multi-view datasets.

Systematically Differentiating Parametric Discontinuities | SIGGRAPH 2021 (ACM ToG.)

S. P. Bangaru*, J. Michel*, K. Mu, G. Bernstein, TM. Li, J. Ragan-Kelley

Link: *https://people.csail.mit.edu/sbangaru/projects/teg-2021/*

- Presented as a talk at *SIGGRAPH 2021* held virtually (originally scheduled to be held at Los Angeles, CA, USA).
- This work systematically addresses the challenges of differentiating parametric discontinuities under integration, by developing a programming language, proving correctness of its semantics, and demonstrating applications including shader design and contact simulation.

Unbiased Warped-Area Sampling for Differentiable Rendering | SIGGRAPH Asia 2020 (ACM ToG.)

S. P. Bangaru, TM. Li, F. Durand   Link: *https://people.csail.mit.edu/sbangaru/projects/was-2020/*

- Presented as a talk at *SIGGRAPH Asia 2020* held virtually (originally scheduled to be held at Daegu, South Korea).
- First paper to show an unbiased area-sampling estimator for the differentiable rendering equation. Derived the estimator from first principles by applying the divergence theorem to the difficult-to-sample boundary term.

Towards Reflectometry from Inter-reflections | ICCP 2020 (Conf.)

K. Shem-Tov*, S. P. Bangaru*, A. Levin, I. Gkioulekas   Link: *https://imaging.cs.cmu.edu/multibounce_reflectometry/*

- Presented as a talk at *ICCP 2020* held virtually (originally scheduled to be held at St. Louis, MO, USA).
- Showcased a differentiable variant of the BDPT rendering algorithm and used it to extract 4-6 times higher fidelity reflectance information by taking advantage of higher-order illumination in concave objects.

Exploration for Multi-task Reinforcement Learning with Deep Generative Models | NeurIPS DRL 2016

S. P. Bangaru, J.S. Suhas, B. Ravindran   Link: *https://arxiv.org/abs/1611.09894*

- Presented as a poster at the *NIPS Deep Reinforcement Learning Workshop 2016* at Barcelona.
- Used a modified version of the deep *variational autoencoder* to encode the various task environments, as the agent explored them. This model is then used by the agent to obtain a Bayesian estimate of the unknown environment, leading to better performance when compared to a standard RL agent.

Talks & Posters

Differentiable Rendering Tutorial | CVPR 2021

Organizers: Prof. Shuang Zhao, Prof. Ioannis Gkioulekas   Link: *https://www.diff-render.org/*

- Virtually presented warped-area sampling and systematic auto-diff for discontinuities at CVPR 2021 to demonstrate advances in differentiable rendering and its applications for computer vision.

DiffSlang: Fast & Reusable Differentiable Graphics Programming with Slang | GTC 2023

Organizers: NVIDIA GTC (GPU Technology Conference)   Link: *Conference Session Catalog | GTC 2023 | NVIDIA*

- Differentiable Slang is an open-source high-performance shader compiler with first-class support for automatic differentiation designed to build and maintain reusable components of large-scale differentiable & neural graphics pipelines.

Awards & Honors

- SIGGRAPH Asia 2023 Best Paper Award (for *path-space warped-area reparameterization*) - ACM
- NVIDIA Graduate Research Fellowship 2022-2023 - NVIDIA
- Meta Graduate Research Fellowship Finalist[1] 2022 - Meta
- Harold Edgerton Fellowship 2019-2020 - MIT EECS
- "C Sivaram Murthy" Best BTech (Undergraduate) Project in Computer Science Award 2017 - IIT Madras
- Charpak Research Scholarship 2016 - French Embassy in India

---

[1] offer declined due to conflict-of-interest

Experience

Research Scientist (Remote) | NVIDIA Research, Cambridge, MA, USA | APR 2024 - CURRENT (FULL-TIME)

Manager: Aaron Lefohn

· Developed (and continuing to develop) novel systems & compilers as a part of the Real-Time Rendering lab, to further the efforts of bringing learning-based components into real-time graphics pipelines.
· Presently focused on developing the differentiable programming features of the Slang shading language.

Research Intern (Remote) | NVIDIA Research, Redmond, WA, USA |
MAY-AUG 2022 & 2023, AUG 2022-MAR 2024 (PART-TIME)

Mentors: Yong He, Tess Foley, Aaron Lefohn

· Worked on *Differentiable Slang*, a strongly-typed shader compiler for high-performance differentiable graphics applications. Implemented a differentiable type system that uses Slang's generics to enable retargetable (reusable) code, as well as an IR-level linearization pass for forward-mode automatic differentiation.
· Can be extended to support several modern differentiable and inverse graphics applications, potentially a large segment of NVIDIA's future outlook.

Research Intern (Remote) | Adobe Research, San Jose, CA, USA | MAY-AUG 2021

Mentors: Kalyan Sunkavalli, Milos Hasan, Michael Gharbi, Zexiang Xu, Sai Bi

· Applied the reparameterization approach to differentiate the sphere tracing process for SDF rendering. Obtained initial results that validate the gradients reparameterized pipeline when compared to ground truth gradients. This is the first pipeline to produce cheap, accurate gradients without blurring/approximating the SDF
· Potentially generalizable to other domains with a compilers approach, in the same way that *Teg* generalized edge-sampling. Project ongoing as a collaboration.

Research Intern | GraphDeco research group, INRIA Sophia-Antipolis, France | MAY-JUL 2017

Advisor: George Drettakis

· Under the supervision of *Dr. George Drettakis*, our team extended the paper *Scalable InsideOut Rendering* to handle thousands of input viewpoints using a streaming architecture and a spatial data structure.
· Played a primary role in developing a framework comprising PatchMatch stereo for mesh refinement, image segmentation for enforcing object boundaries, mesh decimation to control level-of-detail based on distance, and preemptive loading/unloading of images.

Research Intern | GraphDeco research group, INRIA Sophia-Antipolis, France | MAY-JUL 2016

Advisor: George Drettakis

· Worked under the mentorship of *Dr. George Drettakis*, to fix and enhance the pipeline of the popular algorithm *Multi-view intrinsic images with an Application to Relighting* which deals with decomposing the various components of lighting of an outdoor scene using images taken at different parts of the day.
· The project involved a deep understanding of photorealistic rendering and the use of advanced C++ concepts to modify the popular *mitsuba* path tracer, which was used to quantify the algorithm's error at each stage.

Theses

High-fidelity Shape Reconstruction with Differentiable Path tracing | Master's Thesis | JUL-DEC 2018

Advisors: Prof. Ioannis Gkioulekas, Prof. Anat Levin   Link: https://people.csail.mit.edu/sbangaru/diff-pt.html

· Built a differentiable BDPT renderer, that is used with a Tensorflow optimizer to recover precise shape and reflectance properties of the surface, in the presence of complex light transport effects like interreflections.
· This is the first optimization system that can reconstruct generalized (non-convex) meshes with arbitrary surface properties (BRDF) and under arbitrary lighting conditions.

Action-Conditional Projection Neural Networks | Bachelor's Thesis | JAN-APR 2017

Advisor: Prof. Balaraman Ravindran   Link: https://people.csail.mit.edu/sbangaru/acpnn.html

- Applied the idea of ray-tracing to neural networks and redesigned the popular convolutional network with a completely different architecture to simulate and understand 2D/3D interactive scenes.
- Presented reconstruction results with both 2D and 3D perspective projection and provided a compelling argument for ACPNN's advantage over traditional recurrent CNNs.
- Won the *"C Sivaram Murthy" Best Undergraduate Project* award presented by IIT-M

Teaching Experience

6.8300/1 Advances in Computer Vision | Graduate Teaching Staff | MIT EECS | Feb-May 2023

Instructors: Profs. William Freeman, Mina Konakovic & Vincent Sitzmann

Link: http://6.8300.csail.mit.edu/sp23/

- Worked as part of the teaching staff for one of the largest courses in the EECS department (over 150+ graduate and undergraduate students). Helped to run tutorials, held weekly office hours and answered questions on Piazza.
- Reworked an assignment (Problem Set 6) from scratch, making it an interactive Jupyter notebook that can be auto-graded.

Peer-Reviewing Work
- IEEE TVCG 2023 Tertiary Reviewer
- ACM SIGGRAPH 2022 Tertiary Reviewer
- ACM SIGGRAPH Asia 2022 Tertiary Reviewer
- ACM SIGGRAPH 2023 Tertiary Reviewer
- ACM SIGGRAPH Asia 2023 Tertiary Reviewer
- ACM SIGGRAPH 2024 Tertiary Reviewer
- ACM SIGGRAPH Asia 2024 Tertiary Reviewer

Test Scores
- GRE: 332 (Quantitative: 170, Verbal: 162, AWA: 4)
- TOEFL: 112 (Reading: 30, Listening: 30, Speaking: 24, Writing: 28)
- Joint Entrance Examination 2013 Advanced - All India Rank 211 (out of 150,000 shortlisted examinees)
- Joint Entrance Examination 2013 Mains - All India Rank 10 (out of 1.3 million examinees)