Francis Williams

francis@fwilliams.info
+1-650-701-7891

Summary

I am a researcher working at the intersection of computer vision, machine learning, and computer graphics. My work is a mix of theory and application, aiming to solve practical problems in elegant ways. In particular, Im very interested in designing 3D shape representations to enable the use of deep learning on real-world geometric datasets which are often noisy, unlabeled, and consisting of very large inputs. Beyond research, I have a strong background in software engineering and systems programming backed by several years of industry experience and open source projects.

Education

PhD in Computer Science - GPA: 3.98/4.00
New York University, New York, NY
Advisors: Joan Bruna, Denis Zorin
Graduated Suma Cum Laude September 2021

Bachelor of Software Engineering - Honors Co-operative program - GPA: 79.19%
University of Waterloo, ON, Canada
Graduated with Distinction
October 1^{st} 2015

Publications

Neural Fields as Learnable Kernels for 3D Reconstruction - *CVPR 2022*
Francis Williams, Zan Gojcic, Sameh Khamis, Denis Zorin, Joan Bruna, Sanja Fidler, Or Litany

Learning Smooth Neural Functions via Lipschitz Regularization - *SIGGRAPH 2022*
Hsueh-Ti Derek Liu, **Francis Williams**, Alec Jacobson, Sanja Fidler, Or Litany

Neural Splines: Fitting 3D Surfaces with Infinitely-Wide Neural Networks - *CVPR 2021 (Oral)*
Francis Williams, Matthew Trager, Joan Bruna, Denis Zorin

VoronoiNet: General Functional Approximators with Local Support - *CVPR 2020 Workshop*
Francis Williams, Jérôme Parent-Lévesque, Derek Nowrouzezahrai, Daniele Panozzo, Kwang Moo Yi, Andrea Tagliasacchi

Gradient Dynamics of Shallow Univariate ReLU Networks - *NeurIPS 2019*
Francis Williams, Matthew Trager, Claudio Silva, Daniele Panozzo, Denis Zorin, Joan Bruna

Deep Geometric Prior for Surface Reconstruction - *CVPR 2019*
Francis Williams, Teseo Schneider, Claudio Silva, Denis Zorin, Joan Bruna, Daniele Panozzo

ABC: A Big CAD Model Dataset For Geometric Deep Learning - *CVPR 2019*
Sebastian Koch, Albert Matveev, Zhongshi Jiang, **Francis Williams**, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, Daniele Panozzo

Unwind: Interactive Fish Straightening - *CHI 2020*
Francis Williams, Alexander Bock, Harish Doraiswamy, Cassandra Donatelli, Adam Summers, Daniele Panozzo, Claudio Silva

Open Source

Point Cloud Utils
https://github.com/fwilliams/point-cloud-utils

- A Python utility library exposing common algorithms on 3D point clouds and meshes:
  - Random point sampling of meshes with Poisson Disk Sampling and Lloyd-Relaxation
  - Fast pairwise nearest neighbor
  - Point set distances including Chamfer, Sinkhorn, and Hausdorff

NumpyEigen
https://github.com/fwilliams/numpyeigen

- A library for fast zero-overhead bindings between Numpy and Eigen:
  - Makes it easy to transparently convert NumPy dense and sparse arrays into Eigen types while taking full advantage of expression template optimizations in Eigen
  - Features near-zero performance overhead and supports function overloading
  - Used by LibIGL, a major open source project, for Python bindings

FML - Francis' Machine-Learnin' Library
https://github.com/fwilliams/fml

- A collection of Pytorch utilities for machine learning tasks:
  - Includes a numerically stable implementation of the Sinkhorn algorithm to compute optimal transport of point sets in any dimension
  - Also includes a vectorized implementation of the Chamfer-Distance between point sets in any dimension

Unwind
https://github.com/fwilliams/unwind

- A software tool for segmenting and unwarping volumetric CT scans of fishes:
  - Currently deployed in 2 Labs (at Tufts and the University of Washington) with plans for expansion
  - Paper in submission to IEEE TVCG 2019

LibIGL
https://github.com/fwilliams/libigl

- Actively contribute to LibIGL, an open source geometry processing library:
  - Designed, wrote, and maintained new Python bindings (release planned for July)
  - Implemented techniques for meshing
  - Implemented volume rendering in the viewer

Work Experience

Research Scientist Feb. 2021 - Present
NVIDIA (Simulation Research Team) New York, NY

- Conducting fundamental research in deep learning on on 3D shape representations for shape reconstruction and completion of large scale data.

Student Researcher Jun. - Dec. 2020
Google Brain Toronto, ON / New York, NY

- Conducted fundamental research in deep learning on problem of 3D human pose estimation from images
- Developed novel deep learning method for predicting human keypoints in 3D as spatial probability distribution
- Published technical report on the topic to be submitted to a future workshop

Software Engineer Jan 2016. - Jan. 2017
MemSQL San Francisco, CA

- Developed Role-Based Access Control into the MemSQL database engine
- As part of a team of two, designed and implemented Pipelines a high throughput distributed data ingest and transformation engine
- Designed and implemented a subprocess management system in the database engine allowing MemSQL to execute external code in a secure manner without leaking resources
- Designed and implemented a general database backup system which allowed users to specify a backup target (e.g. Amazon S3) and automatically write backups and snapshots to this location

Research Intern Sept. - Dec. 2014
HP Labs, Systems Software Group Palo Alto, CA

- Researched designed and implemented distributed file system to run on a simulated memristor computer with ∼10 TiB of DRAM and hundreds of CPUs with the goal of evaluating performance of file workloads in persistent memory computing environment
- Designed and implemented directories, file descriptor management, and journaling
- Delivered functional prototype with demo which compiled source code of the file system stored in the file system

Research Assistant Jan. 2014 - Apr. 2014
Oregon State University Corvallis, OR

- In collaboration with Dr. Eugene Zhang, researched and developed a correct and efficient algorithm to compute photorealistic lighting in 3D scenes containing multiple interacting mirror surfaces
- Formulated theory to model lighting conditions due to mirrors by understanding how mirrors change the topology of the underlying path space of a scene
- Generated results demonstrating the quality of our method compared to existing techniques
- Designed and developed algorithm as a modification of *POV-Ray*, an existing open source ray tracer
- Independently learned about many new subjects to conduct research, including group theory, topology, and photorealistic rendering

Graphics Software Engineering Intern May - Aug. 2013
Amazon (A9.com) Palo Alto, CA

- Independently developed cross platform graphics framework for rendering 3D scans of products
- The framework supported model and image based rendering techniques with the ability to easily add new rendering modes
- The framework could be deployed to Linux, Windows, Android, iOS and any WebGL compatible browser from a single C++ code-base using *emscripten*