

---

## Education

- 2018–Present **Stanford University.**  
PhD in Computer Science  
Advised by Kayvon Fatahalian
- 2015–2018 **Carnegie Mellon University.**  
MS in Computer Science (Admitted as PhD)  
Advised by Kayvon Fatahalian
- 2011–2015 **The University of California, Davis.**  
BS in Computer Science  
Advised by John D. Owens and Zhendong Su

---

## Major research projects

- 2019–Present **Forager: Rapid Data Exploration and Model Development.**  
It is now possible to create Computer Vision (CV) models that can reliably perform important visual recognition tasks – object detection, classification, depth estimation, etc. However, the techniques for creating such CV models are ad-hoc, expensive, and must be repeated whenever the task definition changes. For example, creating a model for detecting a new type of object requires repeating the manual process of data collection, data labeling, model selection, model training, model optimization, and validation. The Forager project is a broad effort toward unifying and streamlining the model creation process.
- 2015–2018 **Scanner: Efficient Video Analysis at Scale.**  
Scanner is a system for efficient video analysis at scale and one of the major thrusts of the Intel Science and Technology Center for Visual Cloud Systems at CMU (<http://istc-vcs.cmu.edu/>). The focus of Scanner is to answer the question: what are the fundamental primitives for expressing large-scale video processing algorithms and what is the system architecture which can efficiently implement these primitives on heterogeneous clusters featuring both CPUs and GPUs? Scanner is available as an open-source system to the community (<https://github.com/scanner-research/scanner>). In cooperation with collaborators at Facebook, CMU, Stanford, and UC Berkeley, Scanner has been used to:
- Provide the production compute engine for executing Facebook's processing pipelines for synthesizing high-quality 360 stereo VR video.
  - Mine and annotate 10 years worth of TV News (200k hours) to analyze trends and biases in the media
  - Accelerate 3D pose estimation from 480 cameras

---

## Publications

### *Preprints*

- 2021 **Learning Rare Category Classifiers on a Tight Labeling Budget**, Ravi Teja Mullapudi, *Fait Poms*, Bill Mark, Deva Ranman, Kayvon Fatahalian, To appear on Arxiv.

- 2021 **Low-Shot Validation: Active Importance Sampling for Estimating Classifier Performance on Rare Categories**, *Fait Poms\**, *Vishnu Sarukkai\**, *Bill Mark*, *Deva Ranman*, *Kayvon Fatahalian*, To appear on Arxiv.  
*Conference Papers*
- 2021 **Background Splitting: Finding Rare Classes in a Sea of Background**, *Fait Poms\**, *Ravi Teja Mullapudi\**, *Bill Mark*, *Deva Ranman*, *Kayvon Fatahalian*, To appear at CVPR 2021.
- 2018 **Scanner: Efficient Video Analysis at Scale**, *Fait Poms*, *Will Crichton*, *Pat Hanrahan*, *Kayvon Fatahalian*, SIGGRAPH 2018.
- 2018 **Learning Patch Reconstructability for Accelerating Multi-View Stereo**, *Fait Poms*, *Chenglei Wu*, *Shoou-I Yu*, *Yaser Sheikh*, CVPR 2018.

## Talks

*Scanner: Efficient Video Analysis at Scale*

- Strata Data Conference; Moscone West, San Francisco, California; March 2019
- SIGGRAPH 2018; Vancouver Convention Center, Vancouver, Canada; August 2018
- QInF Finals; Qualcomm, San Diego, California; April 2017

*Balancing Systems & Optimization Concerns: Scaling a distributed graph form convex optimization solver with GPUs*

- MOPTA 2015; Lehigh University, Bethlehem, Pennsylvania; July 2015

## Research Experience

- 2015–Present **Ph.D. Student with Professor Kayvon Fatahalian**, CARNEGIE MELLON & STANFORD UNIVERSITY, Pittsburgh, PA & Stanford, CA.  
High-efficiency programmable systems for video analysis
- Nov. 2017 – June 2019 **Research Consultant with Facebook, Surround360 Team**, FACEBOOK, Menlo Park, CA.  
Worked with the Surround360 team at Facebook, who develop omnidirectional stereo cameras for producing feature film quality 360 stereo video for VR, to convert their entire production pipeline to use Scanner to be deployed to film studios for use.
- May – Nov., 2017 **Research Intern with Dr. Shoou-I Yu, Dr. Chenglei Wu**, OCULUS RESEARCH, Pittsburgh, PA.  
  - Submission to CVPR 2018 on learning to predict depth estimation quality from an image for faster and more accurate multi-view stereo reconstruction.
  - Performed analysis and optimization to improve the speed of a GPU-based depth estimation algorithm by 4x over the existing, significantly optimized implementation.
- May – Aug., 2016 **Research Intern with Dr. Michael Garland, Dr. Michael Bauer**, NVIDIA RESEARCH, Santa Clara, CA.  
Large-scale video analysis and Supercomputing runtime systems
- 2012–2013 **Undergraduate Research Assistant with Professor John D. Owens**, UNIVERSITY OF CALIFORNIA, DAVIS, Davis, CA.  
Multi-node Multi-GPU Convex Optimization

2012–2013 **Undergraduate Research Assistant with Professor Zhendong Su**, UNIVERSITY OF CALIFORNIA, DAVIS, Davis, CA.  
Low-level image processing

---

## Industry Experience

2013–2014 **Co-founder**, NIBBOL, San Francisco, CA.

Founded a startup to help the ailing veterinary health-care system. Our focus was on developing software to streamline practice management and analysis and also to make the client's experience more efficient through information transparency.

Detailed achievements:

- Led engineering efforts from prototyping to deployment; the server ran across a cluster of 7 machines on Amazon's EC2, communicating with 35 active client installations in 15 hospitals.
- Developed a soft real-time, multi-user, CRDT-based system for simultaneous and offline modification of client data via a browser.
- Designed and implemented SQL query and data model abstractions that ensured correctness and significantly decreased code footprint.
- Pitched over 280 veterinarians and signed up 55 at platform launch; gave technical explanations and helped 70 front desk receptionists to integrate the software into their workflow.

Summer 2013 **Software Development Intern in Test**, *Microsoft*, Redmond, WA.

Summer 2012 **Graphics Software Validation Engineer Intern**, *Intel*, Folsom, CA.