Favyen Bastani

favyen@csail.mit.edu https://favyen.com/

Education

Ph.D. in Computer Science, Massachusetts Institute of Technology September 2016 – *August 2021*

Thesis: Scalable Video Analytics through Approximate Inference Optimizations

Advisor: Prof. Samuel Madden

B.S. in Computer Science, Massachusetts Institute of Technology September 2012 – June 2016

Publications

- **Favyen Bastani**, Songtao He, Oscar Moll, Sam Madden. SkyhookML: A Platform for Computer Vision Applications. In submission. Demonstration. Website: https://skyhookml.org/.

  SkyhookML is a web-based platform for computer vision, providing an interface geared towards novice users for developing machine learning pipelines involving image and video data. To minimize the effort and time needed to develop such pipelines, SkyhookML incorporates much of our previous research in optimizing video analytics queries, self-supervised learning, and feature extraction from aerial and satellite images. SkyhookML provides tools and optimizations for: (1) annotating images, video, and satellite imagery for image classification, object detection, and image segmentation; (2) training models such as UNet, YOLOv5, as well as architectures from our previous work such as RoadTracer, Sat2Graph, and Self-Supervised MOT; (3) applying these models efficiently on new and large-scale datasets, leveraging our previous work such as MultiScope and Miris; and (4) combining machine learning operations with pre- and post-processing operations to achieve an end-to-end objective.

  We have successfully deployed SkyhookML in several ongoing collaborations. For example, researchers in sociology at Northeastern University used SkyhookML to detect oil wells and pad sites as part of a study on the environmental impacts of drilling activities. We are also exploring applications of SkyhookML in traffic analytics with the Arizona DOT, and in studying the interaction between humans and autonomous vehicles with researchers at the MIT AgeLab.

- **Favyen Bastani**, Sam Madden. MultiScope: Efficient Video Pre-processing for Exploratory Video Analytics. In submission. Access at https://arxiv.org/abs/2103.14695.

  MultiScope efficiently executes queries over video data: for example, a traffic analyst may want to count traffic volumes at hundreds of junctions across a city based on traffic camera video, while an ecologist may want to identify various bird behaviors in thousands of hours of bird feeder video. Prior work in optimizing video analytics queries generally optimize execution speed in only one dimension, ignoring other potential avenues for accelerating execution. In MultiScope, we propose two novel optimization techniques, segmentation proxy models and recurrent reduced-rate tracking, and integrate these and other methods in a cohesive system that tunes parameters across multiple components to provide a superior speed-accuracy tradeoff.

- **Favyen Bastani**, Songtao He, Sam Madden. Unsupervised Multi-Object Tracking with Dual-Tracker Consistency. In submission. Access at https://favyen.com/self-supervised-mot.pdf.

  We propose a novel method, which we call dual-tracker consistency, for training a recurrent multi-object-tracking neural network architecture without video-level annotations. Remarkably, our unsupervised approach provides accuracy that is competitive with three recent fully supervised methods, by training only on a large corpus of unlabeled video.

- **Favyen Bastani**, Songtao He, Satvat Jagwani, Edward Park et al. Updating and Maintaining Street Maps using Changes Detected in Satellite Imagery. In submission. Access at https://favyen.com/detectchange.pdf.

  In this work, we show that existing street map datasets can be automatically and robustly updated with newly constructed roads and buildings by identifying change between two satellite images taken of the same locations at different times. We propose two novel methods, change-seeking iterative tracing and self-supervised selective change detection, that together achieve high-accuracy map update.

- **Favyen Bastani**, Oscar Moll, Sam Madden. Vaas: Video Analytics at Scale. ***VLDB***, 2020. Demonstration. Website: https://vaas.csail.mit.edu.

  Vaas, a predecessor to SkyhookML, is a video analytics system that provides an interactive interface for analyzing large-scale video datasets. Users express workflows as data flow graphs that compose various built-in or custom operations using a query composition tool, and can efficiently experiment with and iterate on these workflows using an interactive execution tool.

- **Favyen Bastani**, Songtao He, Arjun Balasingam, Karthik Gopalakrishnan et al. Miris: Fast Object Track Queries in Video. ***SIGMOD***, 2020.

  Video analytics has become crucial to many applications, including traffic planning and autonomous vehicle development. Miris optimizes the execution of object track queries, providing speedups as high as 17x over prior work, thereby enabling interactive exploration of video at scale.

- Songtao He, **Favyen Bastani**, Satvat Jagwani, Mohammad Alizadeh, Hari Balakrishnan et al. Sat2Graph: Road Graph Extraction through Graph-Tensor Encoding. ***ECCV***, 2020.

  Sat2Graph is a computer vision method for extracting road networks from aerial and satellite images. It recognizes that, broadly, prior methods either apply graph-based tracing (first proposed in our work, RoadTracer) or image segmentation to detect roads. Thus, in Sat2Graph, we propose a novel graph-tensor encoding scheme that combines these two types of methods into a unified framework, thereby deriving a substantial improvement in accuracy, especially for vertically stacked roads such as overpasses, tunnels, and highway interchanges.

- Songtao He, **Favyen Bastani**, Arjun Balasingam, Karthik Gopalakrishnan, Ziwen Jiang et al. BeeCluster: Drone Orchestration via Predictive Optimization. ***MobiSys***, 2020.

  Motivated by the growing number of applications of small aerial drones (e.g., geographic mapping, air pollution monitoring, and search and rescue), BeeCluster proposes a drone orchestration system that manages a fleet of drones on behalf of an application. It provides a virtual drone abstraction so that developers need only express their geographical sensing tasks – BeeCluster automatically determines how to perform these tasks on a fleet of drones. To minimize expected execution time, it employs a novel predictive optimization technique, in which an inferred model of future tasks is used to generate an optimized flight and sensing schedule.

- Songtao He, **Favyen Bastani**, Satvat Jagwani, Edward Park, Sofiane Abbar et al. RoadTagger: Robust Road Attribute Inference with Graph Neural Networks. ***AAAI***, 2020.

  While RoadTracer and other road extraction methods are able to accurately infer road networks, digital street maps include much more about roads than just their positions: road attributes such as the type (residential street vs motorway) and the number of lanes are also crucial. RoadTagger accurately infers these attributes by applying a neural network architecture that combines a convolutional neural network with a graph neural network.

- **Favyen Bastani**, Songtao He, Sofiane Abbar, Mohammad Alizadeh, Hari Balakrishnan, Sanjay Chawla, Sam Madden. Machine-Assisted Map Editing. ***SIGSPATIAL***, 2018.

  Mapping road networks today is labor-intensive. We developed a machine-assisted map editor, MAiD, to enable OpenStreetMap contributors to leverage the outputs of machine learning models trained to detect roads and buildings, thereby improving their map editing productivity. We also propose improvements to the RoadTracer neural network architecture and training procedure to substantially improve its runtime, making it suitable for on-demand execution in MAiD.

- **Favyen Bastani**, Songtao He, Sofiane Abbar, Mohammad Alizadeh, Hari Balakrishnan et al. RoadTracer: Automatic Extraction of Road Networks from Aerial Images. ***CVPR***, 2018.

  RoadTracer is a computer vision method for extracting road networks from aerial and satellite images. Prior work in automatically inferring street maps generally first segments imagery, and then post-processes the segmentation output to extract a road network graph. In RoadTracer, we instead develop a CNN-guided iterative search process that derives a road network directly from the output of the CNN. We show that this approach substantially reduces the error rate.

  RoadTracer has had a large impact on work in the computer vision community studying aerial image processing: a wide range of recent methods have adopted and extended the iterative tracing technique we proposed in RoadTracer, including PolyMapper [ICCV 2019], Neural Turtle Graphics [ICCV 2019], and VecRoad [CVPR 2020].

- Songtao He, **Favyen Bastani**, Satvat Jagwani, Edward Park, Sofiane Abbar et al. RoadRunner: Improving the Precision of Road Network Inference from GPS Trajectories. ***SIGSPATIAL***, 2018.

  RoadRunner is a method for extracting high-precision road networks from GPS trajectory data. Finding that existing techniques suffer from low precision on the very roads with the most GPS trajectories (where high-accuracy should be most achievable), RoadRunner proposes a novel tracing procedure that incrementally follows the flow of trajectories to correctly infer road segment connectivity. RoadRunner not only substantially improves precision over prior work, but it also accurately infers maps even in portions of the road network with complex topologies such as highway interchanges and parallel roads.

- Yan Huang, **Favyen Bastani**, Ruoming Jin, Xiaoyang Sean Wang. Large Scale Real-Time Ridesharing with Service Guarantee on Road Networks. ***VLDB***, 2014.

- **Favyen Bastani**, Yan Huang, Xing Xie, Jason Powell. A Greener Transportation Mode: Flexible Routes Discovery from GPS Trajectory Data. ***SIGSPATIAL***, 2011.

Teaching and Service

October 2020	Reviewer, AAAI 2021
September 2020	Volunteer, VLDB 2020
Spring 2020	Teaching Assistant, 6.824 Distributed Systems
May 2019	Reviewer, IEEE Transactions on Services Computing
August 2018	Reviewer, IEEE Transactions on Services Computing
Spring 2017	Teaching Assistant, 6.S062 Mobile and Sensor Computing