

# Jeet Kanjani

Website LinkedIn Github

jkanjani@andrew.cmu.edu

+1 412-626-1947

## EDUCATION

---

- **Carnegie Mellon University - School of Computer Science** Pittsburgh, PA  
*Master of Science in Computer Vision (4.22/4.33)* Jan 2021 - May 2022
- **The LNM Institute of Information Technology** Jaipur, India  
*Bachelor of Engineering in Computer Science (7.53/10.0)* Jun 2014 - Apr 2018

## ACADEMIC PROJECT

---

- **Far3Det: Towards Far-Field 3D Detection (Argo AI)** Pittsburgh, PA  
*Research Collaborator (Capstone Project)* Jan 2021 - May 2022
  - **Research:** Derived an initial dataset for the problem of far-field 3D object detection (**Far3Det**) and conducted empirical evaluation of baseline algorithms on it. Shown that various late-fusion strategies for multimodal detection applied wrt class specific distance ranges can produce significant improvements on nuScenes dataset. [Link]
- **Robust Factor Graph Attention Net** Pittsburgh, PA  
*(MMML Course Project)* Sept 2021 - Dec 2021
  - **Research:** Conducted experiments to mitigate uncertainty in Visual Question Answering models and make them robust to linguistic variations. Formulated contrastive loss and generated unimodal/multimodal relevance scores for training. Improved on the NDCG metric by **7.3%** over the FGA baseline (Awarded best poster presentation). [Poster][Paper]

## EXPERIENCE

---

- **NVIDIA** Santa Clara, CA (Remote)  
*Software Intern* May 2021 - Aug 2021
  - **Development:** Developed and integrated a modular skeletal based action recognition pipeline into Metropolis MDX (intended for horizontal scalability) and eMDX (intended for low latency edge devices). [Link]
  - Worked with Apache Spark, Apache Kafka, Redis, NVIDIA DeepStream, GStreamer, GRPC, Docker, Triton Inference Server and ELK stack for developing stateful microservices for streaming data.
  - Achieved **2.5X** speedup on Triton Inference Server by building custom op in tensorRT and GPU memory optimizations.
  - Used DetectNet.v2(Resnet34), NvDCF tracker, OpenPose and ST-GCN models in different modules of the pipeline.
  - Containerized all modules of the pipeline and created a python evaluation package with debugging tools for the pipeline.
- **Oxehealth Ltd** Oxford, UK  
*Research Engineer* May 2018 - Dec. 2020
  - **Research:** Contributed towards creating ML models for Person tracking, Fallen Person Detection, Person on Edge of Bed detection and Sleep Staging from video thereby improving the reduction of patients falling from 33% to **48%**. [Link]
  - Achieved a 22% improvement in YOLOv3 mAP by joint learning it with an optical flow based motion model.
  - **Development:** Developed and maintained GRPC services to serve deep learning models in production. Achieved **32%** speedup and deployment cost by 10X by building and shipping inference service with Coral TPU.
  - **Industrialize:** Supervised 3 interns to build active learning annotation tools to reduce time to production. Built an evaluation framework in PySpark for regression testing which is extensively used across teams.
- **University of Oxford** Oxford, UK  
*Research Intern* May 2017 - Sept. 2017
  - **Research:** Worked with a DPhil student at **Torr Vision Group** on 3D Pose Estimation from Monocular images using structured learning approaches. Improved on previously built 2D Pose Estimator using CRF as RNN. [Link]
  - **Development (QuickHOG):** CUDA implementation of HOG-SVM based Pedestrian Detection to the **OxSight glasses** used by the visually impaired. Achieved **80X** run time improvement over sequential implementation and **1.2X** over state of the art parallel implementation. Implemented a novel NMS by adopting a *map/reduce* parallelization pattern. [Link]
- **Tonbo Imaging** Bangalore, India  
*Research Intern* Jan. 2018 - April 2018
  - **Research:** Addressed the issue of long term tracking of objects in thermal infrared videos by using fully convolutional siamese networks (SiameseFC) with LSTMs. Achieved a 2.9% AUC improvement on Tonbo's infrared dataset. [Link]

## PUBLICATION

---

- **Far3Det Towards Far-Field 3D Detection**  
Shubham Gupta\*, Jeet Kanjani\*, Shu Kong, Martin Li, James Hayes, Francesco Ferroni, Deva Ramanan  
*Submitted to CVPR 2021*

## KEY PROJECTS

---

- **ML Toolbox:** pip installable toolbox with standard ML and RL algorithms written using numpy Jan 2021 - Apr 2021
  - Supports Decision Trees, Naive Bayes, Neural Network, Q learning, etc. with evaluation metrics and unit tests [Link]

## TECHNICAL SKILLS

---

- **Programming Languages:** : C, Python, C++, Java, Scala
- **Tools:** Pytorch, JAX, Tensorflow, Caffe, Spark, Pandas, Kafka, Redis, AWS, Docker, SQL, Unity3D, Flask