

# Param Papat

linkedin.com/in/parampopat

ML Engineer — Agentic AI, Multimodal LLM Reasoning, Robotics Simulation

(440) 454-4561

param.papat@gmail.com

## SUMMARY

---

Agentic AI researcher and systems builder with deep experience in autonomous agents, LLM-grounded task planning, multimodal reasoning, and real-world simulation. Currently developing Shortcut Agents at Apple—personalized automation agents that generate and refine Apple Shortcuts through multi-turn reasoning, feedback incorporation, and on-device context. Proven track record of shipping LLM-backed systems, photorealistic 3D simulators, and on-device ML pipelines. Passionate about building general-purpose, safe, and aligned AI systems that integrate language, vision, memory, perception, and control.

## EXPERIENCE

---

- **Apple Inc., Cupertino, CA** *Jan 2021 – Present*  
*Machine Learning Engineer – Agentic AI & Simulation (Special Projects Group)*
  - Led the end-to-end development of Apple’s first agentic LLM interface, enabling users to describe complex tasks and receive executable Apple Shortcuts with over >92% task success across 100+ automation templates.
  - Drove LoRA-based fine-tuning of Apple Foundation Models, integrated 500+ first-party tools, and aligned user experience to deliver a personalized, on-device Shortcut Agent with tool use for adaptive task automation.
  - Integrated multimodal understanding of on-screen context including UI structure, visual semantics, and dynamic user state into a multi-turn Agentic workflow with memory, enabling context-grounded reasoning and iterative refinement with 40% higher user engagement.
  - Led interactive reinforcement learning agents effort capable of exhibiting long-tail adversarial (interactive) behaviors such as tailgating, abrupt braking, etc. that dynamically resized to mimic diverse traffic actors like bikes, pedestrians, and trucks creating diverse training and evaluation scenarios.
  - Developed photorealistic simulation infrastructure using 3D Gaussian Splatting to build 1000+ richly detailed indoor scenes including Human Splats for training and evaluating robotics agents; leveraged super-resolution and multi-GPU training parallelism to scale to 50k+ agent-hours.
  - Constructed city-scale NeRF and 3DGS based digital twins for closed-loop training and evaluation of autonomous systems, enabling sensor-in-the-loop testing, calibration, and perception robustness across occlusion and varying lighting conditions.*AI/ML Software Engineer (Rotation Program)*
  - Built advanced pruning pipelines for neural networks deployed on Apple Silicon, achieving 2x inference speedup without accuracy loss.
  - Co-developed a transformer-based on-device foundation model for Apple Watch gestures, leveraging IMU and PPG signals to power features like Double Tap and AssistiveTouch. The model runs in real-time with minimal battery usage and generalizes across users without requiring explicit calibration.
- **AI Zwei (formerly Dailight)** *May 2020 – Dec 2020*  
*AI Research Intern* *New York, NY*
  - Built transformer based language blocks for production AutoML.
  - Applied distillation, pruning, and quantization pipelines to boost inference speed and robustness.
- **Robert Bosch GmbH** *Jan 2019 – May 2019*  
*Project Trainee (AI)* *Bengaluru, IN*
  - Invented and patented methods to detect adversarial attacks and prevent model stealing; commercialized in Bosch’s AI security service.
  - Built a generative recommendation engine for hypermarket retail, raising conversion rate from 25% to 70%.
- **SnapFactory (Founder)** *Jun 2016 – Jul 2019*  
*AI-Powered Photography Workflow Startup* *Ahmedabad, IN*
  - Built a GAN based low light image enhancer for indoor photo workflows, cutting post-processing time by 60%.

- **Canary Mail** *May 2018 – Jul 2018*  
*Rajkot, IN*  
*Machine Learning Intern*
  - Built GRU based stock seasonality models for NSE equities. Identified buy/sell windows using temporal cross-correlation, achieving simulated 13.5% YoY return.

## EDUCATION

---

- **Columbia University** *2020*  
*New York, NY*  
*M.S. in Computer Science – GPA: 3.98/4.00, TA: Computer Vision, NLP*
- **Nirma University** *2019*  
*Ahmedabad, IN*  
*B.Tech in Computer Engineering – GPA: 9.12/10.00, TA: Deep Learning*

## SKILLS

---

- **Languages & ML:** Python, PyTorch, TensorFlow, Scikit-learn
- **Agentic AI:** LLMs, LoRA, Prompt Engineering, Memory Augmented Planning, Multi-turn Dialog
- **Simulation & 3D:** NeRFs, Gaussian Splatting, Open3D, MeshLab
- **Model Optimization:** Pruning, Quantization, On-device Inference (Apple Silicon)
- **Infra & Tools:** AWS, Git, SageMaker, CloudFormation

## PROJECTS

---

- **Apple Shortcut Agent – LLM based Personal Task Planner** *2024–Present*
  - End-to-end agentic task planner integrated with iOS. Combines LoRA-modified foundation models with UI context ingestion and multi-turn feedback loops.
  - Enabled scalable natural language automation across productivity, utility, and accessibility domains.
- **Are Transformers Learning or Memorizing?** *2020*
  - Trained ALBERT on grammar-constrained data. Findings showed equivalence with lookup tables, exposing LMs' reliance on memorization vs. abstraction.
- **Quick-ML: Email Based ML Prototyping via AWS** *2020*
  - Built serverless ML deployment tool. Users send email to trigger data ingestion, model training, and endpoint provisioning via Lambda, SES, and SageMaker.
- **CNN for Next-Day Stock Prediction** *2019*
  - Used 2D time-series CNN on mapped stock indicators to predict directional movement with F1 = 0.80.

## PUBLICATIONS & IP

---

- **Animal/Object Identification Using Deep Learning on Raspberry Pi** *2019*
  - Published in Springer Smart Innovation Series (with Nirma University).
- **Method to Protect Neural Networks Against Model Extraction** *2019*
  - Patent filed in US, India, Germany – App No: 201941024967.
  - Inventors: Param Popat, Manojkumar Parmar, Himajit Aithal.