Kunsh Singh

bca2nk@virginia.edu | 703-725-7665 | linkedin.com/in/kunsh-singh | github.com/kunshksingh

www.kunshsingh.com

Education

University of Virginia — Computer Science B.S. August 2021—May 2024

Relevant Coursework: Natural Language Processing (Grad), Probabilistic Machine Learning (Grad), Machine Learning (Grad), Machine Learning in Image Analysis, Advanced Software Development, Stochastic Processes, Probability, Linear Algebra, Algorithms, Computer Architecture, Operating Systems, UX Design

Objective

Passionate computer scientist with strong expertise in machine learning, data science, and computational modeling, seeking research opportunities in neurocomputation, specifically focusing on non-invasive brain-computer interfaces and sensory perception modulation.

Skills

Primary Languages: Python**,** C++, C, C#

Frameworks: Kubeflow, FastAPI, Django, Flask, Next.js, React Native

Main Tools: Git, Docker, Elasticsearch, Confluence, JIRA, Obsidian, AI Assisted Tools (Github Copilot + gpt4.5)

Main Libraries: transformers, spaCy, NLTK, Tensorflow, Keras, PyTorch, PyTorch Lightning, Hydra, numpy, scipy, OpenCV, Pandas, scikit-learn, selenium, kaggle

Other Technologies: react.js, highcharts.js, TailwindCSS, Linux/UNIX, HTML/CSS/JS, Unity3D, pyplot, AWS S3

Work Experience

Leidos, *Data Scientist* July 2024—Present

- Designed a fully function AI detection framework prototype within **9 days** leveraging **Flask** and **TailwindCSS**
- Created a reasoning module for the prototype using **LLM with RAG** from AI detection models, enabling it to visually highlight and explain AI generated artifacts with image overlays, togglable with a textual explanation
- Won Leidos' 2024 Shark Tank Pitch Competition with **"Best Pitch"** for AI model framework, presented to **5 VP judges** at Leidos Innovation Showcase as well as over **100 people** on division meeting, **securing additional funding** with customer for **over 3 months** for new project
- Designed a full-stack data discovery platform performing various calculations to visualize time-series data
- Constructed a **FastAPI** backend for data queries and dynamic real time visualizations **every 30 seconds** with **over 10,000 data points** from an S3 bucket, displayed on a frontend utilizing cesium.js and TailwindCSS
- Migrated an existing application from Streamlit to FastAPI improving scalability by containerizing with Docker, and created custom tailwindCSS + cesium.js UI with additional clock synchronized views
- **Reduced debugging times** for various code issues **by over 90%** on the team leveraging generative AI – in one case, received a Kudos for solving an intern's **multi-week CUDA OutOfMemory error** in **less than 1 day**
- Developed and trained a matrix of **30+ deep-learning PyTorch models**, and conducted multiple experiments comparing various weight statistics by layer to determine model similarity
- Conducted Python experimentation with vision transformer and noising algorithms

UVA Advanced Software Development, *Teaching Assistant* Jan 2024—May 2024

- Held **office hours 3 hrs/wk** for students, graded over 500 homework/small-projects/midterms
- Led 2 cohorts, both scored 93% on their semester-long whistleblower webapp in Django leveraging AWS S3

UVA Ultimate User Interface Lab, *Research Assistant* Aug 2023—May 2024

- Combined discrete and continuous redirected walking algorithms using C# in Unity3D and analyzed data in Python, confirming previous literature and suggesting such combination may reduce simulator sickness
- https://doi.org/10.18130/qn2d-k789 [Capstone papers for B.S.]

Leidos, *Data Science Intern* Jun 2023—Aug 2023

- Developed language classification with **spaCy** and **langdetect**, yielding **96.5% accuracy under 5 minutes** on **70,000 brief paragraphs**, used in development of multilingual NER and text embeddings
- Created **multilingual summarization** for over **500 languages**; application uses langdetect, RTG translation (many to English), choice of BART or t5 English summarization in a Flask application, generates ROGUE metrics with SummEval; full model runs in **30 seconds** for **~2-3 multilingual paragraphs** without any scaling
- Researched mT5, LLaMA2, and gpt-3.5-turbo as possible solutions for summarization
- Wrote detailed documentation with Confluence, project used in future R&D with customer (Mar 2025)

UVA Omni-Reality and Cognition Lab (ORCL), *Research Assistant* Jan 2022—May 2023

- Trained YoloV5 (PyTorch) object-detection model with over **1000 manually annotated images** of a VR simulated pedestrian crosswalk to make roads safer/less stressful to cross (epochs=200, batch=16, res=1080)
- **Analyzed eye-tracking data** with OpenCV/NumPy/Tensorflow/Pandas (Python) to determine time viewing objects of interest, and conducted eye tracking and physical to virtual space mapping for lab data collection

Hackathons

- Planned, devised, and implemented 8 hackathon projects, 5 of which were winning projects

Highlighted Projects

Five Nights at Wall Street— Best Financial Hack

- Gamified stock trading in a fast paced, simulated set of 5 consecutive real trading days based on historic data
- Built full stack with Flask to retrieve and cache 1-min resolution intraday stock data from a random ticker on the S&P 500 with StockVantageAPI, frontend webapp with HTML/CSS/JS/Tailwind + dynamic graph generation with Highcharts.js — integrated with AJAX/jQuery and deployed on Intel Developer Cloud

Eye-Tracking Wheelchair — Best Accessibility Hack

- Engineered **iris-tracking** script using **OpenCV** and **Numpy**, transmitting data over a **TCP socket connection** to an **ESP32 microcontroller** to a custom attachment for motorized wheelchairs to allow precise movement

Acne AI

- Developed and served a Flask backend API into a webapp embedding a unsupervised learning transformer block **DinoV2** into a **PyTorch** deep-learning model to classify acne from images
- Gave a fine-tuned analysis with GPT-3.5 for personalized natural remedies, and other treatment options

AI Therapist: Violet v2

- Made AI therapist model with **emotional recognition** via HumeAPI's facial detection + Python, and ElevenLabs' voice cloning in a Django webapp with a GPT-4 backend, enabling **voice-to-voice AI therapy**

ML Project Experience

VyTrader [In Progress]

- Developing a fully autonomous agentic model that trades XAUUSD based on daily patterns in Python
- Invoke calls to GPT4.5 webapp with simulated human-like interaction using selenium, **cutting runtime costs by 50%**, enabling persistent vectorized memory, and tools inaccessible by API (deep-research/web search)

Brain Scan Reconstruction

- Performed reconstruction task on FLAIR subset in BraTS dataset using *medical-diffusion* open-source project; leveraging PyTorch Lightning and Hydra, trained VQ-GAN and latent diffusion model on 20 epochs on a V100
- Added **BERT text-conditioning block** onto model based on HCP dataset **180 subcortical parcellations of 22 brain regions**, and trained latent diffusion model (PyTorch Video Diffusion) on 20 epochs as well on a V100

Mood Generated Playlist

- Performed **mood classification on roughly 278k songs** with RandomForest, XGBoost, and deep learning
- Extracted feature importance in mood determination, and used KNN to generate 20 song playlist
- Utilized Tensorflow/Keras, Sklearn, Numpy, Scipy, and Pandas

Brain Tumor Classification w/ RandAugment

- Used RandAugment (n=3,m=4) on ResNet50v2 to train **brain tumor classification**: achieved **76% accuracy in less than 5 minutes**

MNIST Classification

- Constructed deep-learning linear classifier for MNIST with a best accuracy of 97.94% using **PyTorch**
- Performed binary classification for 0/1 and 6/8 on the MNIST dataset trained with principal components

Publications

Singh, K., & Molloy, K. (2021). The Changes of Cognition in Teenagers after Playing Video Games. *Journal of Student Research*, *10*(2). https://doi.org/10.47611/jsrhs.v10i2.1748

- Studied cognitive changes in video game players over varying post-game time intervals
- Compared cognition before and after video games at different time intervals of 50 VGPs and NVGPs – key inference: a *positive correlation between ADHD and video game playing* (Journal of Student Research)