

PETER
OTT

ott.peter.r@gmail.com
703-336-7082
Chantilly, VA

OttPeterR.com
github.com/OttPeterR
linkedin.com/in/ottpeterr

About Me

I am an accomplished machine learning engineer driven to expand my technical proficiency and pioneer innovative AI solutions. I strive to play a pivotal role in shaping and delivering transformative systems wherever I am.

Education

MS in CS with Machine Learning Concentration - 2024

George Mason University - College of Engineering and Computing

BS in CS, Honors College - 2017

George Mason University - College of Engineering and Computing

Work Experience

BTI360: Data Scientist / Machine Learning Engineer February 2020 - Present

- Engineered a robust news clustering system with AI driven enrichments. Significant technical contributor from prototype to production.
 - RAG, prompt engineering, embeddings, clustering, Docker, Python, SQL
 - Recognized as "the most advanced use case implementation of LLMs in the world" - MSFT Public Sector CTO - late 2023
- Created a purpose-built embedding model for vector search using a repurposed 8B param LLM and fine tuning on a customer-specific data.
 - LLM2Vec, large scale distributed GPU training
- Demonstrated consistent success across multiple domains: Natural Language Processing, text classification, information consolidation and aggregation, and semantic segmentation.
 - PyTorch, HuggingFace, OpenAI, HuggingFace, SageMaker, GPU Acceleration
- Synthesized and presented critical technical reports encompassing team modeling strategies, customer data holding evaluations, and emerging AI technology forecasts.
- Key player for developing internal AI/ML training.

Next Century (acq. by CACI): Machine Learning Engineer January 2019 - February 2020

- Deep learning data fusion research on DARPA's MediFor initiative, focusing on the detection of deepfakes and other image/video manipulations.
 - TensorFlow and PyTorch for model creation and research

IDT: Intern (2013-2017) and Assoc. SWE (2017-2018)

- Machine Learning powered intelligent storage, sorting, and deletion of image files based on relevancy and task importance in a space constrained system

Other Experience

- Rapid discovery of system vulnerabilities via ML powered fuzz testing
Python, genetic algorithm

Supplementary Courses

- HuggingFace - NLP Course
 - Tokenizers, Fine-tuning, Pipeline, Trainer, Datasets

Hobbies

Open Source Contributions

Baby Buddy - Home Assistant Add on

Track baby related tasks as to not forget when the last feeding or nap was. Baby Buddy was an existing application, my work packaged it into an easy to install Home Assistant Add on. Estimated 600 active users.

Docker, NginX, GitHub pipeline/actions

Home Assistant powered home automation

Various Raspberry Pi's and ODroids, and ESP32s hosting home security, surveillance cameras, network ad-blocking, smart light control, presence detection, file storage, family location sharing and more!