Morgan Park, Ph.D. (fictional example)

Software Engineer - New Grad

email@example.edu | github.com/username | linkedin.com/in/username

EDUCATION

Institute of Technology

Ph.D. Computer Science, May 2025 (fictional)

Thesis: Efficient Attention Mechanisms for Long-Range NLP | Advisor: Prof. Jane Smith (fictional)

State University

B.S. Computer Science & Mathematics, 2020 (fictional)

GPA: 3.92 | Summa Cum Laude

PROJECTS

Sparse Attention for Long-Range Transformers

PyTorch, CUDA, Python, W&B

- Designed novel sparse attention mechanism reducing transformer inference cost by 40% on 8K+ sequences
- Trained and evaluated models on 8x A100 GPUs across 3 benchmark datasets
- Published first-author paper at ACL 2024; presented to 200+ attendees

Efficient Training Pipeline for Large LMs

JAX, TensorFlow, TPU, Python

- Developed training pipeline for 1B-parameter language model on TPU v4 pods
- Reduced training time by 25% through custom gradient accumulation and mixed-precision

Distributed Hyperparameter Optimization

Python, Ray Tune, Kubernetes, AWS

- Built distributed hyperparameter search framework running 500+ trials in parallel
- Reduced model tuning time from 3 days to 4 hours across 50 worker nodes

SKILLS

Languages: Python, C++, CUDA, Rust | ML: PyTorch, JAX, TensorFlow, HuggingFace | Tools: Git, Docker, Slurm, AWS, Linux, LaTeX, Kubernetes