

Dmitri Atkinson

Applied Researcher, AI Foundations

📞 (617) 555-2846

✉️ dmitri.atkinson@example.com

🌐 linkedin.com/example/dmitriatkinson

📍 Somerville, MA 02143

STRENGTHS

- 🔗 **Research Ownership**
Owned long running research agendas from question selection through production decisions, and learned which experiments mattered most.
- 🔗 **Reproducible Experimentation**
Versioned data, configurations, and evaluation protocols made difficult findings easier for peers to inspect and trust.
- 🔗 **Cross Functional Translation**
Explained technical tradeoffs to product and business leaders, helping research findings shape practical priorities.
- ⚙️ **Model Efficiency**
Used quantization, sparsification, and checkpointing studies to clarify capacity, memory, cost, and quality choices.
- 👥 **Research Mentoring**
Coached junior researchers through experimental design and writing, and became a frequent sounding board for hard problems.

SKILLS

Foundation Models Multimodal AI

Vision Language Models

LLM Research NLP PyTorch

Hugging Face Distributed Training

RLHF Self Supervised Learning

Robustness Explainability

Model Quantization

Training Optimization

Reproducible Research

LANGUAGES

English

Native

SUMMARY

Applied AI researcher with eight years of post graduate experience advancing **foundation models**, multimodal learning, language models, and scalable deep learning systems. PhD trained computer scientist who owns research agendas from problem selection and experimental design through large scale training, evaluation, reproducibility, publication, and production integration. Deep hands on work with PyTorch, Hugging Face, Lightning, AWS distributed compute, **self supervised learning**, model robustness, explainability, RLHF, model adaptation, and **training optimization**. Research leadership includes 10B plus parameter model development, runs exceeding 500B tokens, **high volume inference**, reusable platform libraries, and controlled experiments. Published NLP and deep learning work at peer reviewed venues. Partners effectively with researchers, engineers, product managers, and business leaders, turning rigorous findings into reliable customer facing AI capabilities.

EXPERIENCE

Principal Applied Researcher, Multimodal AI

Helix Forge AI 📅 January 2024 - Present 📍 Boston, MA

Owns a multimodal foundation model agenda spanning vision language reasoning, safety guardrails, robustness evaluation, model adaptation, and customer facing **production integration**. Leads researchers and engineers across research planning, large scale training, evaluation, reproducibility, publication, and platform delivery.

- Owned multimodal guardrail research with **PyTorch**, improving safety evaluation for customer facing AI products.
- Directed 10B parameter pretraining with AWS distributed compute, supporting runs exceeding 500B tokens.
- Designed controlled **quantization** studies, reducing memory demands and clarifying production model selection.
- Led RLHF and **instruction tuning** with preference data, strengthening deployment readiness through red team findings.
- Shipped reusable inference libraries, giving product teams consistent safety, quality, and latency validation.
- Mentored researchers and engineers, turning experimental findings into **technical reports** and **conference submissions**.

Senior Applied Research Scientist, Model Optimization

Aperture Neural Systems 📅 July 2021 - December 2023 📍 Cambridge, MA

Led optimization research for language and multimodal models across distributed GPU clusters. Owned experiments spanning training efficiency, memory use, inference cost, **model adaptation**, **reproducibility**, and platform integration for production workflows.

- Led distributed GPU optimization research, balancing throughput, memory, **inference** cost, and model quality.
- Built sparsification experiments in PyTorch, creating clearer tradeoffs for large scale training decisions.
- Implemented quantization and **gradient checkpointing**, improving feasible model capacity within cluster memory limits.
- Reproduced peer reviewed **NLP** methods with versioned datasets, protocols, configurations, and documentation.
- Built supervised fine tuning pipelines, moving domain language models from research into evaluated workflows.
- Contributed **distributed training** code, helping successful prototypes reach production engineering teams.

Applied Research Scientist, Deep Learning

Lattice River Research 📅 August 2019 - June 2021 📍 New York, NY

MY CAREER



- Principal Applied Researcher, Multimodal AI at Helix Forge AI (2.8 Years)
- Senior Applied Research Scientist, Model Optimization at Aperture Neural Systems (2.4 Years)
- Applied Research Scientist, Deep Learning at Lattice River Research (1.8 Years)
- Machine Learning Researcher, Applied Modeling at Beacon Ridge Analytics (1 Years)

- Developed publication ready research for language and vision models, combining self supervised representation learning, **robustness** analysis, explainability, model compression, and production oriented inference work across cloud infrastructure.
- Developed self supervised representation learning, improving model quality across language and vision research.
 - Built **PyTorch** pipelines for large text and image corpora on distributed cloud infrastructure.
 - Created **explainability** diagnostics, exposing distribution shift failures before model selection decisions.
 - Implemented **model compression** and inference optimization for prototypes with practical serving constraints.
 - Converted research prototypes into maintained libraries with tests, documentation, and downstream interfaces.
 - Presented limitations and business implications, helping stakeholders make informed model choices.

Machine Learning Researcher, Applied Modeling

Beacon Ridge Analytics 📅 July 2018 - July 2019 📍 Philadelphia, PA

- Built supervised and self supervised **deep learning** systems for language and event sequence data. Established reproducible research practices, investigated optimization methods, and converted prototypes into tested solution code for deployment needs.
- Built language and event sequence models, establishing supervised and self supervised research baselines.
 - Created reproducible **PyTorch** training workflows, giving researchers consistent evaluation starting points.
 - Investigated **optimizer design** through controlled runs, informing model training choices for applied studies.
 - Developed robustness utilities, supporting production model selection under changing data conditions.
 - Created **explainability** tools, helping teams inspect model behavior before deployment reviews.
 - Converted prototypes into tested solution code, supporting inference volume and operational requirements.

PROJECTS

Multimodal Guardrail Evaluation 📅 2025

Built a controlled evaluation framework for vision language safety behavior, combining red team findings, automated checks, explainability signals, and reproducibility artifacts for production decisions.

Large Scale Model Optimization 📅 2023

Compared quantization, sparsification, gradient checkpointing, and training partitioning methods for large language models across distributed GPU environments.

LEADERSHIP & AWARDS

- Best Paper Honorable Mention, NLP Research Workshop, 2021
- Graduate Research Fellowship, Carnegie Mellon University, 2015
- Research Mentor, Applied AI Reading Group, guiding discussion and practical growth
- Reviewer, NLP and Machine Learning Conference Workshops, assessing applied research quality

EDUCATION

PhD in Computer Science, Machine Learning and Natural Language Processing

Carnegie Mellon University 🎓 GPA: 4.0 📅 2018 📍 Pittsburgh, PA

Coursework: *Deep Learning, Natural Language Processing, Machine Learning Theory, Statistical Learning*

Bachelor's Degree in Computer Engineering

University of Illinois Urbana Champaign 🎓 GPA: 4.0 📅 2013 📍 Urbana, IL

Coursework: *Computer Architecture, Algorithms, Probability, Software Engineering*

CERTIFICATIONS

- AWS Certified Machine Learning Specialty 📅 2022
- AWS Certified Cloud Practitioner 📅 2022

TECHNICAL SKILLS

- **Deep Learning Frameworks:** PyTorch, Lightning, Hugging Face
- **Cloud Computing:** AWS, AWS Ultraclusters, Distributed GPU Clusters
- **Model Development:** Foundation Models, LLMs, Vision Language Models
- **Training Methods:** Self Supervised Learning, Supervised Finetuning, Instruction Tuning
- **Optimization Methods:** Quantization, Model Sparsification, Gradient Checkpointing
- **Distributed Training:** Training Parallelism, Partitioning Design, Large Scale Training
- **Model Adaptation:** Parameter Tuning, Transfer Learning, Dialogue Finetuning
- **Safety Research:** Multimodal Guardrails, RLHF, Red Team Evaluation
- **Evaluation Systems:** Robustness Evaluation, Explainability, Validation
- **Data Engineering:** Text Datasets, Image Datasets, Event Datasets
- **Research Delivery:** Controlled Experiments, Reproducibility, Technical Reports
- **Production AI:** Inference Optimization, Model Compression, Platform Libraries

PROFESSIONAL AFFILIATIONS

- Applied AI Reading Group, Research Mentor
- NLP and Machine Learning Conference Workshops, Peer Reviewer

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

REFERENCES

AVAILABLE ON REQUEST