

# ALEXANDER GURUNG

+1 703-835-1897 | [aag1234@gmail.com](mailto:aag1234@gmail.com) | [github.com/alex-gurung](https://github.com/alex-gurung) | US & UK Dual Citizen

## EDUCATION

---

### University of Edinburgh

*PhD in ILCC: Language Processing, Speech Technology, Information Retrieval, Cognition*  
Supervised by Mirella Lapata, expected graduation September 2026

Edinburgh, UK  
Aug. 2023 – Present

### Georgia Institute of Technology

*Master of Science in Computer Science, Concentration in Machine Learning*  
GPA: 3.75/4.0

Atlanta, GA  
Aug. 2020 – Dec 2022

### Georgia Institute of Technology

*Bachelor of Science in Computer Science, Minor in Linguistics*  
GPA: 3.88/4.0, Highest Honors

Atlanta, GA  
Aug. 2018 – Dec. 2020

## RESEARCH INTERESTS

---

I'm currently working on Reinforcement Learning for long-story generation, and am focusing my thesis on improving reasoning across long texts. I am especially interested in more complex reasoning incorporating implicit knowledge and multiple data sources, as well as long-term planning.

## RESEARCH EXPERIENCE

---

### University of Edinburgh

*Phd Candidate*

Atlanta, GA  
Aug. 2023 – Present

- Working on long-narrative generation and modeling believable characters.
- Currently exploring latent variable modeling for story generation.
- Designed improved character-representation, CHIRON, to ensure character consistency (EMNLP 2024)

### Social and Language Technologies (SALT) Group @ Georgia Tech

*Researcher*

Atlanta, GA  
Jan. 2020 – May 2023

- Investigated the processes and distribution of radicalization on insular social medias
- Worked with the School of International Affairs to create a radical-online-content ontology
- Analysed prevalence of political frames using dependency-parsing system
- Finetuned and domain-adapted language models to detect radical content and assess its distribution across multiple alt-tech platforms
- Identified dehumanization language as a gateway to further radicalization, and quantified its spread amongst far-right recruiters

### Meta AI

*AI Resident - ParLAI Team*

New York City, NY  
Aug. 2021 – Sep. 2022

- Worked with the *LIGHT* team to improve commonsense understanding in text-adventure games
- Imbued language models with an understanding of world-state and ability to predict state changes
- Designed crowdsourcing tasks to collect a large dataset of game playthroughs, action-result pairs, and human evaluations
- Created novel grounding tasks to improve a model's ability to reason about its environment
- Fine-tuned language models and showed improvement over non-grounded baselines on human evaluations

### Electro-Optical Systems Laboratory @ Georgia Tech Research Institute

*Machine Learning Graduate Research Assistant*

Atlanta, GA  
Jan. 2021 – May 2021

- Developed ML, CV, and DSP solutions for the *Electronic Warfare Modeling and Analysis Division*
- Leveraged adversarial neural techniques for data augmentation to improve generalization performance
- Expanded Genetic Programming framework's CV capabilities with image feature extraction techniques

### Automated Algorithm Design Lab @ Georgia Tech

*Undergraduate Researcher*

Atlanta, GA  
Jan. 2019 – May 2020

- Optimized cache invalidation for lab's framework improving results by 213%
- Led new NLP team in adding core text embedding functionality

## PUBLICATIONS

---

### **Learning to Reason for Long-Form Story Generation**

*Alexander Gurung, Mirella Lapata*

*Preprint, 2025*

### **CHIRON: Rich Character Representations in Long-Form Narratives**

*Alexander Gurung, Mirella Lapata*

*EMNLP Findings, 2024*

We design an improved character representation, CHIRON, for downstream story tasks and analysis.

### **Infusing Common-Sense Reasoning Models with Graph Knowledge**

*Alexander Gurung, Jack Urbanek, Arthur Szlam, Jason Weston*

*2023*

Improved LLM performance on text-adventure game tasks by training on synthetic graph data.

## WORK EXPERIENCE

---

### **University of Edinburgh**

*Teaching Assistant @ Foundations of NLP*

*Mountain View, CA  
December 2024 – Present*

- Designing course materials for undergraduate NLP course
- Creating coding and theory assignments for sentiment analysis and translation

### **TikTok**

*Machine Learning Engineer Intern - Trust & Safety Team*

*Mountain View, CA  
Jun. 2021 – Aug. 2021*

- Designed MoE neural architectures to improve region-specific auto-moderation performance
- Applied and built upon research into multi-task learning loss functions and architectures
- Demonstrated improvements in auto-moderation performance over existing models
- Deployed new models to production and evaluated changes in performance over time

## TECHNICAL SKILLS

---

**Programming Languages:** Python, Java, HTML/CSS/JS, Typescript, C, MySQL/SQL, GoLang, Matlab, R, Dart

**Frontend Frameworks:** React, Angular, React-Native, Flutter, Material-UI, Android

**Backend Frameworks:** NodeJS, Flask, LoopBack, SQL/PostgreSQL, GCP, AWS, Firebase, GraphQL

**Data Science/ML:** PyTorch, Tensorflow, Keras, NLTK, Gensim, SciPy, NumPy, Pandas

**Languages:** English (Native), French (Proficient)