

# Neeraja Kirtane

kirtane3@illinois.edu

Website ◊ LinkedIn ◊ GitHub ◊ Google Scholar

---

## RESEARCH INTERESTS

**Research Interests:** Deep Learning, Natural Language Processing, Graph Deep Learning.

**Focus Areas:** Trustworthiness in NLP, NLP for low-resource languages, Class imbalance handling in Graph Learning.

---

## EDUCATION

<b>University of Illinois Urbana-Champaign</b>	2023 – 2025
<i>Masters of Science in Computer Science (Thesis Track) (MSCS)</i>	CGPA: 4.0/4.0
<b>Manipal Institute of Technology, Manipal, India</b>	2018 – 2022
<i>B.Tech in Computer Science and Engineering (Minor: Computational Intelligence)</i>	CGPA: 9.14/10

---

## EXPERIENCE

<b>University of Illinois, Urbana-Champaign</b>	Aug 2023- Present
<i>Graduate Student Researcher</i>	<i>Advisor – Prof. Hao Peng</i>

- Working on various problems to make Large language models (LLMs) more efficient and trustworthy and to mitigate hallucinations in these models.
- Investigating how hidden layers play a role in detecting hallucinations and facts.

<b>Indian Institute of Technology Madras, Chennai, India</b>	Jul 2022- Aug 2023
<i>Post Baccalaureate Fellow</i>	<i>Advisors – Prof. Balaraman Ravindran &amp; Dr. Rajashree Baskaran</i>

- Worked on the [Project Hidden Voices](#) at the **Robert Bosch Centre for Data Science and Artificial Intelligence (RBC-DSAI)**.
- Building intelligent tools to aid in adding 10,000 notable women’s biography drafts to Wikipedia.
- This aims to **reduce the gender gap** in wikipedia data.
- Worked on building knowledge graphs and doing graph to text generation using large language models.
- Finetuned models like GPT-J and GPT-Neo for this process.

<b>Indian Institute of Technology Madras, Chennai, India</b>	Jan 2022- Jun 2022
<i>Research Intern</i>	<i>Advisors – Prof. Balaraman Ravindran &amp; Dr. Ashish Tendulkar</i>

- Worked on handling class imbalance in Graph neural networks at **RBC-DSAI**.
- Used **implicit ways** at the algorithmic level to handle this imbalance.
- Used a custom loss function and tuned the attention weights to focus more on minority nodes.
- Additional Links: [Report](#) | [Slides](#) | [Github](#)

<b>Centre for development of advanced computing, CDAC Pune</b>	Jun 2020 – Aug 2020
<i>ML Intern</i>	<i>Advisor – Rahul Dangi</i>

- Extracted keywords and named entities from a document for **better comprehension**.
- Used word embeddings of the GLoVe dataset for the predictions. Major libraries used in Python were NLTK (for text processing), Gensim (to use the LDA algorithm), Flask (to create the front end of the project).
- Created an **application** so that people could use it.
- Additional Links: [Github](#) | [Report](#)

---

## PUBLICATIONS

### 1. Wikiworkshop 2023

Hidden Voices: Reducing gender data gap, one Wikipedia article at a time [Paper](#) May 2023

- Authors: *Neeraja Kirtane, Anuraag Shankar, Chelsi Jain, Ganesh Katrapati, Raji Baskaran, Balaraman Ravindran*
- Proposed an algorithm to automate writing wiki articles for women in STEM.
- Discussed the challenges and limitations of the problem at hand.

### 2. GCLR workshop at AAAI 2023

ReGrAt: Regularization in graphs using attention mechanism to handle class imbalance [Paper](#) Sep 2022

- Authors: *Neeraja Kirtane, Jeshuren Chelladurai, Balaraman Ravindran, Ashish Tendulkar*
- Devised a **custom loss** function by adding a regularizer that handles imbalance.
- Used attention mechanism by making the attention weights focus more on minority nodes, in node classification.
- Our results outperformed the already existing methods by a margin of **5 %**.

### 3. Deployable-AI workshop at AAAI 2023

Efficient Gender Debiasing of Pre-trained Indic Language Models [Paper](#)

Aug 2022

- Authors: *Neeraja Kirtane, V Manushree, Aditya Kane*
- Created a template-based dataset suitable for Hindi language to measure bias.
- Measured bias in Hindi Language model by predicting mask probability of a noun/pronoun given the occupation in the template created.
- Debaised the model by efficiently finetuning **unfreezing less than 1 %** of the parameters by training on a balanced dataset.

### 4. Gender bias in NLP workshop at NAACL 2022

Mitigating gender stereotypes in Hindi and Marathi [Paper](#)

May 2022

- Authors: *Neeraja Kirtane, Tanvi Anand*
- Created a dataset of **160-plus gendered and neutral occupations** in Hindi and Marathi. Also created a dataset of emotions broadly classified into anger, fear, joy, sadness.
- Proposed methods to quantify the bias in the word embeddings by modifying ECT and RND tests for gendered occupations.
- Defined a gender axis, using Principal Component Analysis (PCA). Neutralized and debaised the embeddings by removing this component from the embeddings.
- Additional Links: [Slides](#) | [Poster](#)

### 5. WASSA workshop at ACL 2022

Transformer based ensemble for emotion detection [GitHub](#) | [Paper](#)

Mar 2022

- Authors: *Aditya Kane, Shantanu Patankar, Sahil Khose, Neeraja Kirtane*
- Developed ensemble based solution consisting of multiple *ELECTRA* and *BERT* models.
- Proposed methods for **synthetically generating datasets** to mitigate class imbalance.
- Additional Links: [Experiments](#) | [Slides](#) | [Poster](#) | [Video](#)

### 6. Widening NLP workshop at EMNLP 2021

Occupational Gender Stereotypes in Indian Languages [Paper](#) | [Video](#) | [Poster](#)

Nov 2021

- Authors: *Neeraja Kirtane, Tanvi Anand*
- Devised a metric similar to WEAT to calculate bias in **gendered languages** like Hindi and Marathi.
- Used this metric on ULMFiT language model and quantified the occupational bias present.

## PROJECTS

---

### Evaluating Mathematical Reasoning Chains [Github](#)

- Developed a pretrained metric to evaluate the chains generated by LLMs for Math reasoning tasks.
- The metric evaluated nine different characteristics of the chain.

### Smart Document Explorer [GitHub](#)

Summer 2021

- Created a program to make a document more **accessible and understandable**. Used history and geography textbooks as the data to help children benefit from this.
- Extracted named entities, keywords. Summarized the text, found similar sentences given a sentence.
- Used relationship extraction to map dates with events in history textbooks.

## TECHNICAL SKILLS AND RELEVANT COURSEWORK

---

**Languages:** Python, C++, Java, C, SQL

**Tools and Libraries:** PyTorch, NumPy, TensorFlow

**Courses:** CS 412: Introduction to Data mining, CS 546: Advanced NLP, CS 562: Advanced Topics in Security, Privacy, and Machine Learning, CS 568: User-centered ML

## TEACHING EXPERIENCE AND EXTRACURRICULAR

---

- **TA for CS 105:** Introduction to Computing: (Non-Tech) for Fall 2023, Spring 2024
- Volunteer at **EMNLP 2021, NAACL 2022**
- **Regional Mathematics Olympiad (RMO)** Finalist.
- College level Finalist at **Smart India Hackathon** among forty plus teams