

RESEARCH INTERESTS

ML for social good & impact via responsible, safe, trustworthy, and ethical ML solutions that improve downstream critical decision-making. Also interested in building end-to-end solutions for measurable impact.

EDUCATION

- **Johns Hopkins University** Baltimore, MD
Masters of Science in Engineering, Computer Science; CGPA: 3.97/4.00 *Aug 2022 – May 2024*
- **Indraprastha Institute of Information Technology, Delhi** New Delhi, India
B. Tech., Computer Science & Applied Mathematics; CGPA: 9.23/10.00 *Aug 2016 – Aug 2020*
 - **Department Gold Medal** (awarded for best academic performance) [\[link\]](#)[\[cert\]](#)
 - **Thesis:** Deep Mean Shift Clustering. [\[link\]](#)

PUBLICATIONS

- White, J., Madaan, P., Shenoy, N., Agnihotri, A., Sharma, M., & Doshi, J. (2022). A Case for Rejection in Low Resource ML Deployment. ArXiv preprint [arXiv:2208.06359](https://arxiv.org/abs/2208.06359). [Poster in *Challenges in Deploying and Monitoring ML Systems Workshop at NeurIPS 2022*] [\[poster\]](#) [\[preprint\]](#)
- Madaan, P., & Sadat, F. (2020, May). Multilingual Neural Machine Translation Involving Indian Languages. In the *5th Workshop on Indian Language Data: Resources and Evaluation (pp. 29-32) at LREC 2020*. [\[link\]](#)

EXPERIENCE

- **Johns Hopkins University** Baltimore, MD
Graduate Research Assistant
 - **Natural Language Text Data Anonymization** *Aug 2023 - Present*
Advised by: Prof. Anjalie Field [Collaboration with Amazon Research Lab]
 - * Developing end-to-end sensitive data detection and anonymization model for long-form unstructured natural language data by modeling it as a translation task.
 - * Building the mechanism geared toward long-form documents found as medical notes, courtroom proceedings, and child welfare data. Using a Longformer model to fit all the tokens in the context window instead of chunking.
 - * Measuring performance with a suite of metrics covering detection, anonymization, and utility.
 - * **Keywords:** *Privacy, Anonymization, Natural Language Text*
 - **Language Agnostic Sentence Representations sans Parallel Data** *May 2023 - Present*
Advised by: Prof. Phillip Koehn
 - * Developing a model capable of producing sentence representation agnostic of the language without parallel data.
 - * The approach models Transformers as AutoEncoders and uses only unpaired data from multiple languages to build representations agnostic of the source language.
 - * Such representations would have the potential to be used in any downstream Natural Language Task without language being the barrier to access.
 - * **Keywords:** *Representation Learning, Sentence Vectors, Monolingual Data, Low Parallel Resource*
 - **Domain Adaptative Malarial Slide Analysis** *Nov 2022 - Aug 2023*
Advised by: Dr. Ben Haeffele & Matthew Ippolito, M.D., Ph.D.
 - * Implemented Mean-Teacher and Gradient Reversal Domain Adaptation methods for Malarial Image Slide Analysis improving performance on new unseen domains.
 - * Used Contrastive Predictive Coding for cell detection along with Maximal Coding Rate Reduction Loss to improve model performance on noisy labels.
 - * **Keywords:** *Domain Adaptation, Object Detection, Convolutional Sparse Coding, Noisy Labels, Healthcare*

Wadhvani Institute for Artificial Intelligence

Mumbai, India (Virtual)

- *Associate Machine Learning Scientist - I*
Research Fellow

Jan 2021 - July 2022

Jul 2020 - Jan 2021

- **Pest Management | CoreML [code] ★ 8 on **

Managers: Dr. Jerome White & Jigar Doshi

- * Developed a flexible & generic Object Detection codebase with rejection, visualization & deployment capabilities, on top of PyTorch Lightning, Hydra, & NNI.
- * Reduced the Mean Absolute Error from 4 to 2 for the existing object detection system, with the addition of new architectures, compression algorithms & a rejection framework to build trust incorporating on-ground feedback.
- * Solution reached thousands of small landholder farmers in three of the highest cotton-producing states of India
- * Solution won the Global Change Award 2022 [article].
- * **Keywords:** *Object Detection, Open Source, Model Pruning, Framework Building, Core ML, Rejection System, Trustworthy, Deep Learning, Model Deployment, Social Good*

- **Université du Québec à Montréal**

Montreal, Canada

- **MITACS Globalink Research Intern**

May 2019 - Aug 2019

- **Multilingual Neural Machine Translation for Low-Resource Languages**

Advised by: Prof. Fatiha Sadat

- * Developed neural multilingual translation model for low-resource languages.
- * Solution in the top 3 in half of the translation tasks in LoResMT SharedTask at MT Summit 2019
- * Improved the BLEU score by 15 points from the baseline.
- * **Keywords:** *Neural Machine Translation, Low Resource, Transformers*

ACADEMIC SERVICE

- **Course Assistant Computation Finance:** Conducted office hours, and graded assignments.
- **Course Assistant Statistical Analysis:** Conducted office hours, and graded assignments.
- **Program Committee - RANLP 2023:** Reviewed papers submitted in DravidianLangTech 2023 Workshop at RANLP 2023.
- **Reviewer - ACM TALLIP 2022:** Reviewed papers submitted in ACM TALLIP 2022.
- **Reviewer - DravidianLangTech-ACL 2022:** Reviewed papers submitted in DravidianLangTech-ACL 2022 Workshop.
- **Course Assistant Gateway Computing Python:** Conducted office hours, labs and in-class help sessions, and graded assignments.


AWARDS

- Part of Wadhvani AI team that won the HUL, Google, and MyGov India's AI for Agriculture Hackathon. The winning prize was a grant of 1 Million INR.
- **Mitacs Globalink Research Internship 2019:** One of few selected students for a fully funded research opportunity at **Université du Québec à Montréal**.
- Best Academic Performance in Major [cert]
- Dean's Academic Excellence Award for 2 consecutive years: 2017-18, 2018-19 [cert]

SKILLS

- **Masters Courses:** Machine Learning System Design, **AI Ethics**, Computer Vision, Network Security, **Machine Translation**, Machine Learning, Causal Inference
- **Bachelors Courses:** Deep Learning[†], Machine Learning, Speech Recognition[†], Affective Computing[†], Reinforcement Learning[†], Linear Algebra, Probability and Statistics, Real Analysis, ODEs & PDEs, Calculus in \mathbb{R}^n [†], Scientific Computing, Numerical PDEs[†], Differential Geometry, Linear Optimisation
([†]Graduate level courses)
- **Tools & Technologies:** Python, Java, C++, PyTorch, HuggingFace, GPT4ALL, LangChain, Pinecone, Jupyter, Git, Torchtext, Streamlit, Numpy, scikit-learn, PyTorch Lightning, Detectron2, Docker, Hydra, NNI.

PROJECTS

- **SageRef: Single Image Reflection Removal** [\[code\]](#) [\[report\]](#)
 - Modeled a denoising autoencoder to remove image reflection from mirror reflection-plagued images with only a single image.
 - Used a UNet model with synthetic data from the SIR2 Benchmark.
 - **Keywords:** *PyTorch Lightning, TorchVision, TorchMetrics*
- **Benoit: Better English Noisy Audio Transcripts** [\[code\]](#) [\[slides\]](#)
 - Developed a grammar-correcting ASR model for non-native English speaker audio.
 - Created synthetic dataset by back-translating English sentences from a low-resource language and passing them to Microsoft SAPI5 TTS to create a proxy for non-native English audio.
 - Used a GRU-based seq2seq denoising autoencoder on top of a pre-trained Wav2Vec 2.0 (frozen) for grammatically correct ASR.
 - **Keywords:** *PyTorch, TorchAudio, TorchText, Colab*
- **LLM Chatbots** [\[code\]](#)
 - Built a chatbot with LangChain prompt templates, HuggingFace google-flan-t5-xl model on vectorized Anime database indexed in Pinecone.
 - Built a chatbot with mistral-7b-instruct from GPT4ALL that replies with shell commands that can achieve the task asked by the user in the prompt and runs inside a shell can can execute those commands if accepted by the user.
 - **Keywords:** *HuggingFace, Transformers, LangChain, PineCone, LLM, Pinecone, GPT4ALL*
- **Emotional Text-to-speech** [\[webpage\]](#) [\[slides\]](#) ★ 311 ♡ 46 on 
 - Developed over Tacotron for emotional speech synthesis for English.
 - Explored fine-tuning approaches for pre-trained models to synthesize emotional speech using ~15 mins. of audio.
 - **Keywords:** *Deep Learning, Speech Synthesis, Tacotron*
- **Flow Based Generative Models: GLOW** [\[code\]](#) [\[slides\]](#)
 - Explored Conditional GLOW in different generation and conversion tasks as a replacement for Vcoders and GANs
 - **Keywords:** *PyTorch, Colab, Flow Models*
- **Doom Playing DeepRL Agent** [\[code\]](#)[\[slides\]](#)
 - Trained an agent using Deep Recurrent Q-Learning to play Doom: An FPS game having partially observable 3D states.
 - Recreated a simple case from the [Arnold](#) framework with a [Deep Recurrent Q-Network](#).
 - Added the capability to self-learn as the agent plays against self to train itself.
 - **Keywords:** *PyTorch, ViZDoom, Reinforcement Learning, Deep Q-Learning*
- **Passive v/s Active Induced Emotions**
 - Comprehensive Analysis on widely used affective features.
 - Analyzed featured importance in predicting affective state in an active v/s passive visuals study.
 - Collected gameplay videos for active visuals and Bollywood movie trailers for passive.
 - Collected affective scores for the curated dataset.
 - **Keywords:** *PyTorch, Librosa, Scikit-learn*

CO-CURRICULAR ACTIVITIES

- Part of the organizing Committee, [ACSS'18](#) (Workshop on AI for Computational Social Systems)
- Research Volunteer with [Global Village Foundation](#). Helped conduct surveys in rural areas and analyzed data on the impact of government policies and schemes.
- Served as the Event Head, RoboWars, ESYA'17 (Technical Fest)