Krishnapriya Vishnubhotla

vkpriya@cs.toronto.edu • https://priya22.github.io/ • NLP | ML

Education

PhD in Computer Science — **Toronto, Canada**
University of Toronto, Computational Linguistics Group, GPA: 3.97/4.0 — *2019–2024*
Supervised by: Graeme Hirst, Frank Rudzicz

M.Sc in Computer Science — **Toronto, Canada**
University of Toronto, Computational Linguistics Group, Thesis option, GPA: 4.0/4.0 — *2017–2019*
Supervised by: Graeme Hirst

B.Tech in Computer Science and Engineering — **Mangalore, India**
National Institute of Technology Karnataka-Surathkal, CGPA 8.93/10 — *2013–2017*

Experience

National Research Council (NRC) Canada — **Research Intern**
Emotion Dynamics | Mental Health | Social Media — *June 2023 – Jan 2024*
Supervised by: Saif M. Mohammad
Developed the Utterance Emotion Dynamics framework to characterize temporal patterns in emotional expression and variation in social media utterances. Identified systemic variations in Emotion Dynamics metrics with demographic characteristics of users, including as an indicator of multiple mental health conditions.

Georgian Partners — **Research Intern**
Unsupervised Representation Learning | Clustering — *May 2020 – Dec 2020*
Worked on unsupervised auto-encoder models of learning aspect-specific entity representations and clustering algorithms for internal company applications.

Samsung AI Research Center — **Research Intern**
Multi-modal Representation Learning | Video-caption alignment — *May 2019 – Sep 2019*
Worked on multi-modal representation learning and semi-supervised methods of text and video alignment. I was a part of the winning submission for the Samsung Retail Robot Challenge, for which we built an interactive clip retrieval system for customer support videos.

Myntra Designs Inc. — **Software Engineering Intern**
Customer Chatbot | Ruby on Rails — *May 2016 – July 2016*
Worked on backend and frontend development of web interfaces for customer service chatbots.

Indian Institute of Technology Bombay — **Summer Research Intern**
Game Theory | Operations Research — *May 2015 – July 2015*
Worked on characterizing Nash equilibrium of quasi-zero-sum games.

Research Projects

Comparing Lexical Biases in Character Portrayal in Human and LLM Story Generations
Created the GPT-WritingPrompts dataset, which pairs short stories written by Reddit users in response to prompts with comparable generations from GPT-3.5. Quantified and compared biases in character portrayal when grouped by the generation process, narrative voice, and the gender of the protagonist along six lexical dimensions, and showed that the distribution of these features can distinguish machine-generations from human-written stories.

Quantifying Emotional Variation in Temporal Utterance Sequences
Developed the framework of Utterance Emotion Dynamics to characterize temporal patterns of emotion variation in textual utterances in multiple social media datasets. Demonstrated systemic variation in UED metrics based on the geographic and demographic characteristics of speakers. In collaboration with Affective Science researchers, developed a metric to capture Emotion Granularity from textual utterances, which functions as a reliable indicator of multiple mental health conditions.

Improving Automatic Quotation Attribution in Literary Novels
Created the Project Dialogism Novel Corpus (PDNC), the largest current dataset of characters and their attributed quotations in long-form, full-length literary novels. Benchmarked state-of-the-art pretrained models for named entity recognition, coreference resolution, and quotation attributions on the dataset. Developed BERT-based models to attribute quotations to the speaking character entity, which yielded a $> 10\%$ improvement over state-of-the-art.

Characterizing Semantic Relatedness of Sentence Pairs
Created a dataset of sentence pairs annotated for real-valued scores of semantic relatedness using comparative annotations and best-worst scaling. Quantified the contribution of various linguistic features to relatedness of sentence pairs, and finetuned BERT-based sentence representation models to improve state-of-the-art on predicting semantic relatedness. In collaboration with multiple researchers, expanded this to a SemEval-2024 shared task on multilingual Semantic Textual Relatedness.

- **Learning Disentanged Representations of Content and Style in Texts**
  Evaluated autoencoder-based neural seq2seq models to learn disentangled representations of content and style on multiple style transfer datasets. Conducted detailed ablations on a highly-structed Natural Language Generation dataset to identify the effects of motivational, adversarial, and contrastive losses in learning effective disentangled representations.

- **A Stylometric Investigation of Character Voices in Literary Fiction**
  Developed metrics to quantify variation in stylistic and emotional characteristics of character dialogue in a corpus of modern-era plays, and a corpus of full-length fiction novels. Demonstrated systemic biases in certain lexical features of character utterances based on gender of the speaker and gender of the author with mixed-effect models. Qualitatively investigated connections between metrics of dialogic diversity and authorial style.

- **Learning Dicrete Latent Structure: GANs for Text Generation with word2vec intermediaries**
  Developed a text generation model using Generative Adversarial Networks (GANs), with word2vec embeddings as the input and output of the system. This bypasses the problem of differentiating through a discrete space, i.e, words. Our method is agnostic to vocabulary size and achieves competitive results relative to methods with various discrete gradient estimators.

Publications

- Vishnubhotla, K., Teodorescu, D., Feldman, M.J., Lindquist, K.A. and Mohammad, S.M., 2024. Emotion Granularity from Text: An Aggregate-Level Indicator of Mental Health. To Appear, *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (EMNLP 2024)*

- Huang, X.Y., Vishnubhotla, K. and Rudzicz, F., 2024. The GPT-WritingPrompts Dataset: A Comparative Analysis of Character Portrayal in Short Stories. To Appear, *Proceedings of the 6th Workshop on Narrative Understanding (WNU-2024)*

- Vishnubhotla, K., Hammond, A., Hirst, G., and Mohammad, S.M., 2024. The Emotion Dynamics of Literary Novels. In *Findings of the Association for Computational Linguistics (ACL 2024)*

- Ousidhoum, N., Muhammad, S.H., Abdalla, M., et al., 2024. SemRel2024: A Collection of Semantic Textual Relatedness Datasets for 14 Languages. In *Findings of the Association for Computational Linguistics (ACL 2024)*.

- Vishnubhotla, K., Rudzicz, F., Hirst, G., and Hammond, A., 2023. Improving Quotation Attribution in Literary Novels. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers) (ACL 2023)*

- Abdalla M., Vishnubhotla, K. and Mohammad, S.M., 2023. What Makes Sentences Semantically Related: A Textual Relatedness Dataset and Empirical Study. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (EMNLP 2023)*

- Vishnubhotla, K., Hammond, A. and Hirst, G., 2022. The Project Dialogism Novel Corpus: A Dataset for Quotation Attribution in Literary Texts. In *Proceedings of the 13th Language Resources and Evaluation Conference (LREC 2022)*

- Vishnubhotla, K. and Mohammad, S.M., 2022. Tweet Emotion Dynamics: Emotion Word Usage in Tweets from US and Canada. In *Proceedings of the 13th Language Resources and Evaluation Conference (LREC 2022)*

- Vishnubhotla, K., Hirst, G. and Rudzicz, F., 2021. An Evaluation of Disentangled Representation Learning for Texts. In *Findings of the Association for Computational Linguistics (ACL 2021)*

- Vishnubhotla, K., Hammond, A. and Hirst, G., 2019. Are Fictional Voices Distinguishable? Classifying Character Voices in Modern Drama. In *Proceedings of the 3rd Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2019)*

- Budhkar, A., Vishnubhotla, K., Hossain, S. and Rudzicz, F., 2019. Generative Adversarial Networks for Text Using Word2vec Intermediaries. In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*

Technical skills

- **Deep Learning Frameworks:** PyTorch, TensorFlow, JAX, AWS Cloud and SageMaker
- **ML/NLP Libraries:** HuggingFace, LangChain, spaCy, NLTK, Pandas, NumPy, Scikit-learn
- **Programming Languages:** Python, R, Javascript, C, C++, Java

Teaching Experience

- **Co-Instructor, University of Toronto:** Introductory Computation and Data Science for the Life and Physical Sciences, *(Winter 2024)*

- **Lead Teaching Assistant, University of Toronto:** Introductory Computation and Data Science for the Social Sciences, Life and Physical Sciences, and Literature, *(Winter 2023, Fall 2023)*
- **Teaching Assistant, AI4Good Lab:** The AI4Good Lab is a 7-week introductory machine learning program for under-represented groups in the field. As a TA, I designed and ran daily tutorial sessions on ML theory, programming, and mentored teams on an applied NLP project on style transfer. *(Summer 2022)*
- **Teaching Assistant, University of Toronto:** Natural Language Computing, Introduction To Computer Programming, Introduction to Computer Science, Programming on the Web. *(2017 to 2022, various semesters.)*

Academic Service

- Research Mentorship: Xi Yu Huang, May-December 2023.
- I have served as a reviewer for:
  - ACL Rolling Review: 2024, 2023, 2022
  - *ACL Conferences (ACL, NAACL, EMNLP): 2022, 2021, 2020, 2019, 2018
- Volunteer mentor for the Graduate Application Assistance Program, 2021.
- Triager for the DCS Admissions Program, 2021 and 2020.
- Maintained the official webpage of the Computational Linguistics group, University of Toronto, from 2018-2020.

Awards

- Vector Research Grant, Vector Institute, Toronto *(2019–2024)*
- SGS and DCS Conference Grants, University of Toronto *(2021, 2023)*
- Ministry of Human Resources and Development India Undergraduate Scholarship *(2013–2017)*

Relevant Courses

(Advanced) Computational Linguistics ○ Natural Language Computing ○ Introduction to Machine Learning ○ Learning Discrete Latent Structure ○ Algorithms for Private Data Analysis ○ Topics in Computational Social Science ○ Computational Models of Semantic Change