



CS 333: NLP

Fall 2025



Prof. Carolyn Anderson
Wellesley College

Final Project Presentations

- ◆ 2-3 minutes per person
- ◆ 1 slide (upload to Google Drive by 10pm on Monday)



CS DEPARTMENT PRESENTS

winter Wonderland

DECEMBER

11

4:30 PM - 6 PM

AT THE LEAKY BEAKER AREA



BRACELET MAKING,
MEXICAN HOT CHOCOLATE,
MULLED CIDER & COOKIES



QUESTIONS? SB129 ACCOMMODATION? ACCESSIBILITY@WELLESLEY.EDU



PLAYABLE MEDIA SHOWCASE

THURSDAY,
12/11 @ 2-4 PM
IN SCI-L140

?: ad114@wellesley.edu

ADR: accessibility@wellesley.edu

Help hours

My help hours for the rest of the semester:

- ◆ **Next Monday:** 4-5
- ◆ **Next Thursday:** 3:30-4:30
- ◆ **Following Monday:** 4-5

What Have We Learned?

Tasks We Learned This Semester

- ◆ Tokenization
- ◆ Sentence Segmentation
- ◆ Textual Similarity Analysis
- ◆ Representation Learning
- ◆ Classification: sentiment analysis, genre prediction...
- ◆ Language Modeling
- ◆ Common-sense reasoning

Techniques We Learned This Semester

- ◆ **Tokenization**: BPE
- ◆ **Textual Similarity Analysis**: tf-idf, PPMI
- ◆ **Representation Learning**: term-document frequency matrixes, static word embeddings, contextualized word embeddings
- ◆ **Classification**: Naive Bayes, regression, neural networks
- ◆ **Language Modeling**: n-grams, neural networks (feedforward, recurrent, attention-based)
- ◆ **Transfer Learning**: fine-tuning, prompt engineering, proxy-tuning
- ◆ **Mechanistic Interpretability**: activation steering

Models We Used This Semester

- ◆ word2vec
- ◆ regression and feedforward neural networks in PyTorch
- ◆ Encoder-decoder Transformers (DistilBERT)
- ◆ Decoder-only Transformers (Llama, Smol)

Skills We Developed This Semester

- ◆ **text processing**: working with strings, reading in text datasets, tokenizing, cleaning text, regular expressions
- ◆ **evaluation**: computing metrics (accuracy, precision, recall, F-score, perplexity, and cosine similarity), understanding statistical reliability
- ◆ **implementing classification models**
- ◆ **building** and **running** neural network models
- ◆ **fine-tuning** and **prompting** large language models

What Is There Still
To Learn?

Interpretability and Explainability

How do blackbox models make decisions?

What are large language models learning, and how do they use this information?

What are the limits of what large language models?

How can we intervene when models are incorrect?



BlackboxNLP 2023

The sixth edition of BlackboxNLP will be co-located in Singapore on December 7 2023.

News

- We welcome EMNLP Findings papers with interpretability for presentation at the BlackboxNLP workshop. Contact us via blackboxnlp@googlegroup.com by October, noon ET.
- Find us on Twitter/X here: <https://twitter.com/blackboxnlp>.
- We started a YouTube channel: <https://youtube.com/@blackboxnlp>. Subscribe to be informed about the latest research in NLP!

Workshop on analyzing and interpreting neural networks for NLP



TrustNLP: Third Workshop on Trustworthy Natural Language Processing

Colocated with the Annual Conference of the Association for Computational Linguistics (ACL 2023)

About

Recent advances in Natural Language Processing, and the emergence of pretrained Large Language Models (LLM) specifically, have made NLP systems omnipresent in our daily lives. While these systems have revolutionized many aspects of our society, they also raise important questions about their trustworthiness, fairness, and accountability. The TrustNLP workshop aims to bring together researchers, practitioners, and stakeholders from various fields to discuss these issues and develop new methods and tools for building trustworthy NLP systems.

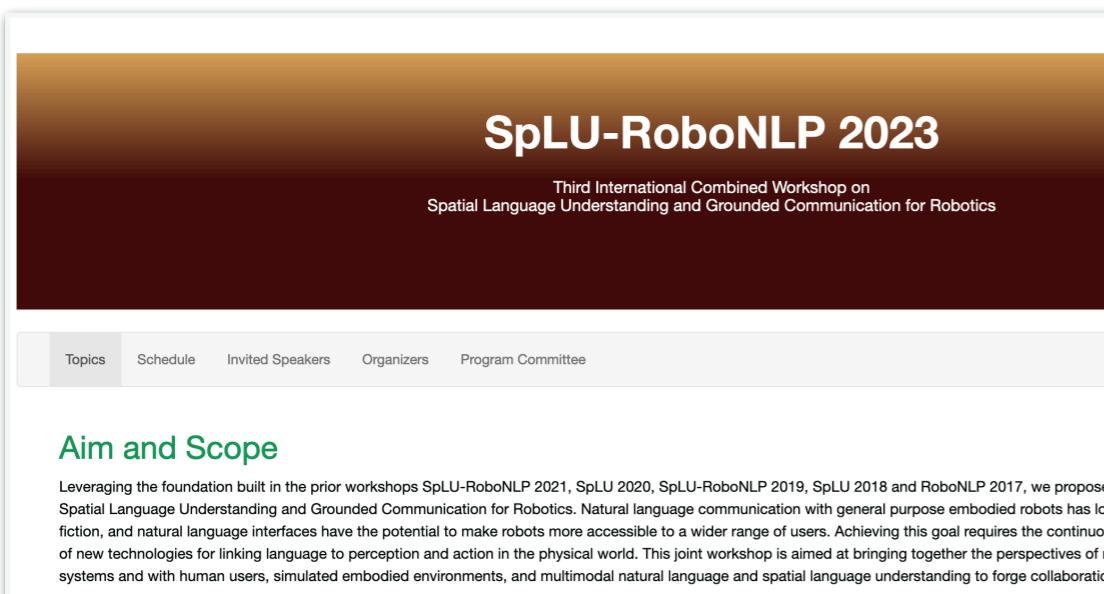
Grounding and Multi-Modality

How do we combine text-based models with other modalities, like vision and perceptual information?

How can we help models relate to the physical world?

How do we help models understand social interaction?

How much of a role does access to perceptual information play in learning?



SpLU-RoboNLP 2023
Third International Combined Workshop on
Spatial Language Understanding and Grounded Communication for Robotics

Topics Schedule Invited Speakers Organizers Program Committee

Aim and Scope
Leveraging the foundation built in the prior workshops SpLU-RoboNLP 2021, SpLU 2020, SpLU-RoboNLP 2019, SpLU 2018 and RoboNLP 2017, we propose the Spatial Language Understanding and Grounded Communication for Robotics. Natural language communication with general purpose embodied robots has long been a goal of robotics, and natural language interfaces have the potential to make robots more accessible to a wider range of users. Achieving this goal requires the continuous development of new technologies for linking language to perception and action in the physical world. This joint workshop is aimed at bringing together the perspectives of robotics systems and with human users, simulated embodied environments, and multimodal natural language and spatial language understanding to forge collaboration.

Confirmed Speakers Call For Papers Submission Instructions Schedule Papers Organizers

NILLI
Novel Ideas in Learning-to-Learn through Interaction
[All Editions](#)

- [NILLI First Edition @ EMNLP'21](#)
- [NILLI Second Edition @ EMNLP'22](#)
- [NILLI Second Edition @ EMNLP'23](#)

Workshop @ EMNLP 2023
Collaborative dialogues [1, 2, 3] with automated systems through language interactions have become ubiquitous, wherein it is becoming common from setting an alarm to planning one's day through language

Linguistic Diversity

How can we improve NLP tools for low-resource languages?

How can we achieve good performance with sparse data?
With small models and limited compute?

How do models deal with different language varieties?

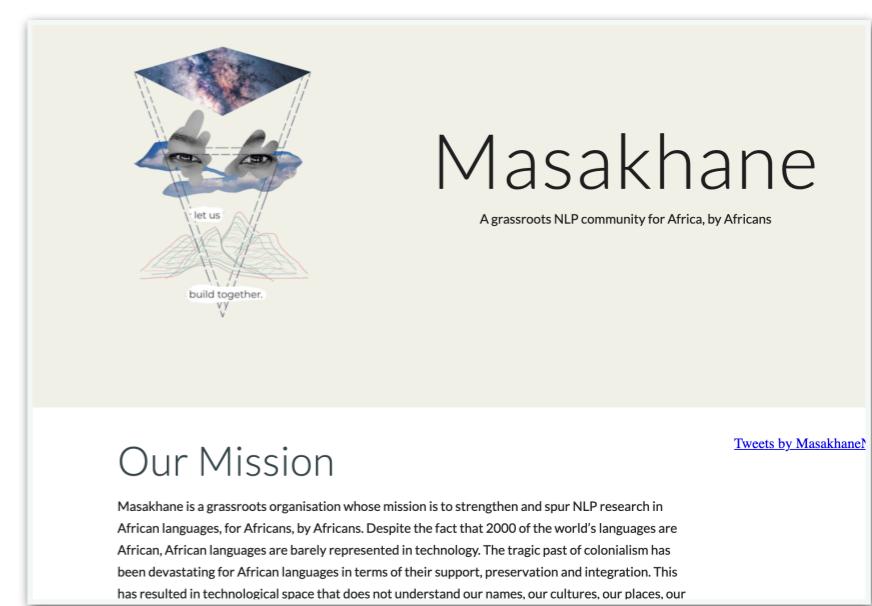


The screenshot shows the AmericasNLP website. At the top, there is a decorative header with a green and red sun-like graphic. Below it, a navigation bar with links: 'AmericasNLP Workshop', 'AmericasNLP 2024 Shared Task', 'AmericasNLP 2023 Workshop', 'AmericasNLP 2023 Shared Task', and '2022 Co...'. The main content area features the text 'AmericasNLP' in large, bold, dark green letters. Below this, a sub-section titled 'AmericasNLP aims to...' lists the following goals:

- ...encourage research on NLP, computational linguistics speech around the globe to work on native American languages
- ...connect researchers and professionals from underrepresented speakers of endangered languages with the machine learning community



The screenshot shows the SustaiNLP 2023 workshop website. The header features the text 'COMPUTEL WORKSHOP' and 'The Use of Computational Methods in the Study of Endangered Languages'. Below the header, there are navigation links: 'ComputEL Workshop', 'ComputEL-7 Workshop (2024) ▾', and 'Previous Workshops ▾'. The main title 'SustaiNLP 2023' is prominently displayed, followed by the subtitle 'Fourth Workshop on Simple and Efficient Natural Language Processing'. At the bottom, there is a 'Workshop Description' section.



The screenshot shows the Masakhane website. The header features the text 'Masakhane' and 'A grassroots NLP community for Africa, by Africans'. Below the header, there is a graphic of a hand holding a small plant. The text 'let us build together.' is overlaid on the graphic. In the bottom right corner, there is a link 'Tweets by Masakhane'.

Our Mission

Masakhane is a grassroots organisation whose mission is to strengthen and spur NLP research in African languages, for Africans, by Africans. Despite the fact that 2000 of the world's languages are African, African languages are barely represented in technology. The tragic past of colonialism has been devastating for African languages in terms of their support, preservation and integration. This has resulted in technological space that does not understand our names, our cultures, our places, our

Computational Social Science

How can we share our best approaches for text processing with other fields?

How well do NLP techniques work with limited data?
In niche domains? In tasks that require expert judgment?

SIGHUM
ACL Special Interest Group on Language Technologies for the Socio-Economic Humanities

5th Workshop on Narrative Understanding

LaTeCH-CLfL 2023
The 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature

TADA 2023: New Directions in Analyzing Text as Data

Today's Guest