Language Technologies Institute

Faculty Research Areas

Carnegie Mellon University
School of Computer Science

KEYWORDS AND AREAS OF RESEARCH

NATURAL LANGUAGE PROCESSING

- Natural Language Analysis (Syntax / Semantic / Pragmatic Analysis) (Fahlman, Fried, Levin, Mitamura, Neubig, Rosé)
- Question Answering (Mitamura, Neubig, Nyberg)
- Natural Language Generation (Fahlman, Fried, Ippolito, Li, Mitamura)
- Conversational AI, Intelligent Agents, and Dialogue (Bigham, Bisk, Busso, Fried, Mitamura, Nyberg, Rosé, Rudnicky, Sap, Wantanabe)
- Efficient NLP (Li, Neubig, Strubell)
- Multilingual NLP (Brown, Frederking, Levin, Li, Mitamura, Mortensen, Neubig, Singh, Waibel)
- Machine Translation (Brown, Frederking, Levin, Li, Mortensen, Neubig, Waibel)
- Information Extraction (Frederking, Mitamura, Mortensen, Neubig, Rosé, Strubell)
- Reasoning (Fahlman, Levin, Li, Rosé, Strubell)
- Fairness and Ethics in Language Technology (Diaz, Sap, Strubell)
- Creativity (Diaz, Ippolito)
- Evaluation (Diaz, Li)

COMPUTATIONAL LINGUISTICS

- Morphology and Phonology (Levin, Mortensen)
- Morphosyntax and Syntax (Levin, Mortensen)
- Semantics (Fahlman, Mitamura)
- Discourse and Pragmatics (Frederking, Fried, Rosé, Sap)
- Sociolinguistics (Rosé)
- Language Change (Mortensen)

SPEECH PROCESSING

- Speech Emotion Recognition (Busso)
- Speech Recognition (Raj, Rudnicky, Singh, Stern, Waibel, Watanabe)
- Speech Synthesis (Waibel, Watanabe)
- Multilingual/Low-Resource Speech Processing (Brown, Busso, Frederking, Levin, Mortensen, Singh, Waibel, Watanabe)
- Speech Translation (Frederking, Li, Waibel, Watanabe)
- Speech Forensics (Raj, Singh)
- Speech Enhancement / Robust Speech Processing (Raj, Singh, Stern, Watanabe)

MULTIMODAL LEARNING

- Multimodal AI (Bisk, Busso, Fried, Hauptmann, Mitamura, Morency, Rosé, Rudnicky, Singh, Waibel)
- Multimedia analysis (Hauptmann, Rudnicky, Waibel)
- Language + Vision (Bisk, Fried, Morency, Singh)
- Affective Computing (Busso, Morency, Singh)
- RoboNLP / Embodied AI (Bisk, Fried, Rudnicky)
- Social Signal Processing (Busso)
- Nonverbal behaviors generation (Busso)
- Multimodal applications (Busso)

INFORMATION RETRIEVAL

- Recommender Systems (Callan, Diaz, Nyberg)
- Retrieval and Ranking Models (Callan, Diaz, Nyberg)
- Search, Recommendation, and RAG (Callan, Diaz, Nyberg, Xiong)

MACHINE LEARNING FOR LANGUAGE TECHNOLOGIES

- Graph-based Machine Learning (Yang)
- Neural Network Algorithms (e.g., XL-Net, DARTS, etc.) (Busso, Fahlman, Neubig, Raj, Rosé, Strubell, Waibel, Yang)
- Time Series, Spatiotemporal Modeling (Rosé, Shamos, Yang)
- Query-driven Graph Generation for Causality Analysis (Yang)
- Extreme-scale Text Classification (Yang)
- Language Technology Application Areas/Issues (Bigham, Ippolito, Mitamura, Nyberg, Shamos, Strubell, Yang)
- Privacy and Security (Ippolito, Li, Raj, Shamos)
- Code generation (Fried, Neubig, Li, Welleck)
- Computational Social Science (Rosé, Sap)
- Language Technology in Healthcare and Mental Health (Bigham, Busso, Hauptmann, Morency, Rosé, Singh)
- Language Technology in Education (Mitamura, Rosé)
- Fairness and Ethics in Language Technology (Sap, Strubell)
- Interpretability and Explainability in Language Technology (Rosé)

KNOWLEDGE-BASED AI

- Natural-language understanding and generation (Fahlman, Mitamura)
- Knowledge-driven hierarchical planning (Fahlman, Mitamura)

FOUNDATION MODEL CORE

- Efficient Pretraining (Xiong)
- Scaling Up (Xiong)
- Post-Training (Xiong)
- Data-Centric Methodologies (Xiong)

AI FOR SCIENCE AND MEDICINE

- AI for Drug Discovery (Li)
- Healthcare AI (Xiong)
- Protein Modeling and Design (Li)

Jeff Bigham I Associate Professor

HUMAN COMPUTER INTERACTION
- Accessibility for People with Disabilities
- Design, Safety & Responsible AI
- Human-Agent Teaming

Yonatan Bisk I Assistant Professor

EMBODIED AI
- Can we learn language from robots?
- Can we control robots with language?

GROUNDING / MULTIMODAL
- Can models capture both concrete and abstract thought?
- What knowledge about the world cannot be learned from text?

Ralf Brown I Principal Systems Scientist

INFORMATION EXTRACTION
- Text normalization

Language Identification
- Identifying short texts in 1000+ languages

DIGITAL FORENSICS
- Text extraction
- Reconstructing corrupted file

Carlos Busso I Professor

SPEECH AND MULTIMODAL SYSTEMS

Affective Computing
- Speech emotion recognition
- Multimodal Machine Learning
- Speech processing
- Social Signal Processing
- Nonverbal behaviors for conversational agents
- Multimodal applications

Jamie Callan I Professor & Director, PhD Program

SEARCH ENGINE ARCHITECTURES
- Discrete and continuous text representations
- Heuristic and neural search architectures
- Search using knowledge and semi-structured information
- Conversational search

THE LEMUR PROJECT
- Open-source search engine and text mining software
- Large-scale web datasets

Mona Diab I Professor | LTI Director | ACL Fellow

COMPUTATIONAL LINGUISTICS/NATURAL LANGUAGE PROCESSING
- Arabic NLP
- Conversational AI
- Computational lexical semantics
- Multilingual and cross lingual processing & Low resource processing
- Computational socio-pragmatics
- Information extraction & Text Analytics
- Machine Translation
- Data Resource creation & Annotation Science
- Evaluation Science

RESPONSIBLE AI
- Automated accountability
- Trustworthy AI
- Bias/Fairness Research
- Automated Regulatory Compliance
- Diversity and Inclusion

Fernando Diaz I Associate Professor

INFORMATION RETRIEVAL ALGORITHMS
- Design of search engines that enhance machine learning systems
- Support of tip of the tongue information needs

EVALUATION OF LANGUAGE TECHNOLOGIES
- Measurement of unfairness and other harms in language technologies
- Identification and measurement of the impact of artificial intelligence on culture industries such as music, film, and literature
- Design of effective and efficient foundational metrics for offline and online evaluation

Scott Fahlman I Professor Emeritus, LTI & CSD

SYMBOLIC KNOWLEDGE REPRESENTATION AND REASONING
- Ongoing research on the open-source Scone knowledge base system
- Flexible, human-like, "good-enough" planning, integrated with world-knowledge

KNOWLEDGE-BASED NATURAL LANGUAGE UNDERSTANDING AND GENERATION
- Going all the way from text or speech to a useful representation of the knowledge
- Using context and background knowledge for disambiguate and fill in missing information

INCREMENTALLY CONSTRUCTED NETWORKS FOR DEEP LEARNING
- Updating some old ideas about gradually building up neural networks to fit the task at hand

Robert Frederking I Assoc. Dean, Doctoral Programs, SCS & Chair, MLT Program

Daniel Fried I Assistant Professor

GROUNDED INTERACTION
- Grounding language to perception and action
- Language interfaces

CODE GENERATION
- Language-to-code
- Interaction with code generation models

APPLIED PRAGMATICS
- Understanding implicit language
- Contextual language understanding & generation

Alex Hauptmann I Research Professor

MULTIMEDIA ANALYSIS AND RETRIEVAL
- Large scale analysis of internet and surveillance video
- Multimedia for healthcare
- Multimedia for human rights

Daphne Ippolito I Assistant Professor

PRIVACY AND SECURITY ISSUES IN LANGUAGE MODEL SYSTEMS
- Memorization
- Data poisoning
- Prompt extraction
- Detection of generated text

EFFECT OF TRAINING DATA ON LANGUAGE MODEL CAPABILITIES
- Impact of data curation choices

USES OF NATURAL LANGUAGE GENERATION
- Building interactive tools for creative writers
- Real-world evaluation of language generation systems

Lori Levin I Research Professor

MULTILINGUAL NATURAL LANGUAGE PROCESSING

- Using linguistics and language typology to improve NLP
- Low resource languages
- Corpus annotation (syntax, semantics, morphosyntax)

NORTH AMERICAN COMPUTATIONAL LINGUISTICS OPEN COMPETITION (NACLO)

- Chair
- Students learn about linguistics and computation by solving puzzles

Lei Li I Assistant Professor

LARGE LANGUAGE MODELS

- Security, Safety, Privacy, Copyright of LLM
- LLM Agent, reasoning, code generation, tool use, cooperation
- Efficient LLM systems

MULTILINGUAL NLP

- Machine Translation (how to translate 1000 languages?)
- Speech Translation
- Multilingual Transfer

AI FOR SCIENCE

- AI for Drug Discovery, design effective small molecule drugs for disease target
- Generative Algorithms for Protein Design (design proteins with desired biochemical functions

Teruko Mitamura I Research Professor & Director, MIIS Program

TEXT ANALYSIS

- Events: Definition, detection, coreference, sequence, linking and representation
- Annotation on event and entity mentions and linking

QUESTION ANSWERING

- Question answering on various domains
- Question generation and answering from text

Louis-Philippe Morency I Associate Professor

ARTIFICIAL SOCIAL INTELLIGENCE

- Analyze, recognize, and predict subtle human communicative behaviors during social interactions Multimodal Machine Learning
- Learning probabilistic and neural models from heterogenous, contingent, and asynchronous data Health Behavior Informatics
- Technologies to support clinical practice during diagnosis and treatment of mental health disorders

David R. Mortensen | Assistant Research Professor

COMPUTATIONAL MODELS OF LANGUAGE STRUCTURE
- Modeling language change and historical reconstruction of languages
- Modeling and exploiting sound structure and word structure in natural language processing
- Linguistic evaluation of large language models

COMPUTATIONAL MODELS OF LANGUAGE DIVERSITY
- Linguistic typology
- Culture variation in natural language processing

LANGUAGE TECHNOLOGIES FOR PEOPLE WITH DISABILITIES

Graham Neubig | Associate Professor

MACHINE LEARNING FOR NATURAL LANGUAGE PROCESSING
- Large Language Models and Applications
- Automated Machine Learning (AutoML)
- Efficient and Effective Evaluation

NATURAL LANGUAGE INTERFACES TO COMPUTERS
- Code Generation
- AI Agents and Chatbots

MULTILINGUAL AND MULTIMODAL LANGUAGE PROCESSING
- Low-resource Language Processing
- Multilingual Multimodal Models
- Computational Linguistics

Eric Nyberg | Professor

OPEN ADVANCEMENT OF QUESTION ANSWERING
- Software architectures and algorithms for real-world QA applications (e.g., Jeopardy! Challenge, BioASQ, LiveQA)

MACHINE READING
- Mixed-initiative information extraction, logical form creation and inference for automatic knowledge base construction in any domain

INTERACTIVE ANALYTIC LEARNING
- Reducing cost of training high-quality analytics for new domains

Kemal Oflazer | Teaching Professor

TEXT SIMPLIFICATION
- Generating short summaries or headlines from (Turkish) news stories

NLP APPLICATIONS
- Paraphrasing short questions in q/a settings

Using NLP techniques to automate and improve student Q/A in Piazza-like courseware settings

COMPUTATIONAL MORPHOLOGY
- Going beyond morphological segmentation for inducing morpho-semantic representations
- Handling segmentation ambiguity

Bhiksha Raj I Professor

AUDIO ANALYSIS
- Audio content analysis, with applications to acoustic intelligence, surveillance, content-based retrieval
- Never ending learner of sound: a self-updating audio-content index for the web
- Signal enhancement and separation algorithms

PRIVACY PRESERVING SIGNAL PROCESSING
- Algorithms to preserve user privacy in speech & audio applications
- Secure cloud computing techniques
- Speech Processing
- Robust speech recognition and core speech recognition technologies

DEEP NEURAL NETWORKS
- Novel applications of deep networks and algorithms for practical deployment of deep networks

Carolyn P. Rosé I Professor LTI & HCII & Director, MCDS Program

TEXT MINING/ COMPUTATIONAL SOCIOLINGUISTICS
- Modeling social processes in discourse
- Deep learning of rhetorical structure
- Social Media Analysis
- Medical NLP

DIALOGUE AGENTS
- Reinforcement Learning for Adaptable Dialogue Agents
- Dialogue agents for Learning, Health, and Wellbeing Computer-Supported

COLLABORATIVE LEARNING
- Architectures for supporting online collaboration
- Social Recommendation Algorithms
- Learning in Massive Open Online Courses

Roni Rosenfeld I Professor

MACHINE LEARNING APPLICATIONS
- Epidemic Tracking and Forecasting (Delphi Group)
- AI and Society

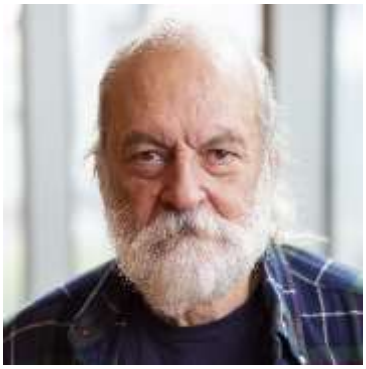

Alexander Rudnicky I Research Professor Emeritus

CONVERSATIONAL AI AND SPEECH RECOGNITION
- Open-Domain Conversation Architectures
- Goal-Oriented Dialog Systems
- Spoken Language Understanding and Situational Awareness
- Emotion Recognition from Speech Audio
- Speech Emotion Recognition and Speech Synthesis

SPOKEN LANGUAGE SYSTEMS
- Dialog System Architectures and Dialog Management Using Implicit State Spaces
- Spoken Language Understanding and Situational Awareness Based on Levels of Context
- Speaker emotion recognition from speech audio

Maarten Sap I Assistant Professor

SOCIAL INTELLIGENCE IN AI AGENTS
- Assessing Theory-of-Mind, and pragmatics of AI systems
- Developing simulation frameworks for LLM agents

MEASURING THE ETHICS AND HARMS OF AI
- Alignment biases in AI (culture, race, gender)
- Frameworks for AI governance
- Human-AI reliance and anthropomorphism

COMPUTATIONAL STORYTELLING ANALYSES
- Commonsense inference with stories
- Empathy and HCI with stories

TOXICITY AND BIAS IN LANGUAGE
- Explainable hate speech detection and counterspeech generation

Michael Shamos I Distinguished Career Professor & Director, MSAII program

MATHEMATICAL DISCOVERY
- Mathematics As a Language
- Automated Discovery of Mathematical Relationships Via AI
- Mathematical Searching

AI LAW
- Development of Meaningful AI Regulations

Rita Singh I Research Professor

SPEECH AND MULTIMODAL SYSTEMS
- Voice Intelligence and Security
- Human Profiling from Voice
- Generalized AI Models for Speech and Audio Processing
- Multimedia and Cyber Forensics
- Human-Guided and Creative AI

QUANTUM COMPUTING

Richard Stern I Professor

SPEECH
- Robust Speech Processing
- Speech Enhancement
- Far-Field Speech Processing
- Auditory and Speech Perception
- Digital Signal Processing

Emma Strubell I Assistant Professor

EFFICIENT NLP/GREEN AI

- How to obtain state-of-the art model accuracy while reducing computation, memory, carbon footprint?
- Which model parameters, training examples are necessary/sufficient for learning effective models?
- How to set up parameter learning, model architecture to facilitate efficient inference?

ROBUST OUT-OF-DOMAIN/OUT-OF-DISTRIBUTION PERFORMANCE

- Transfer learning, learning from few examples, weak supervision.
- How to effectively integrate structured information/priors alongside distributed representations?

PRACTICAL STRUCTURED INTERFACES FOR NATURAL LANGUAGE TEXTS

- Representations that facilitate learning/inference as well as analysis by end-users

Alex Waibel I Professor

- Speech-To-Speech Translation
- Neural Network / Deep Learning and Language Processing
- Machine Learning
- Machine Translation
- Speech Processing
- Multimodal And Multimedia

Shinji Watanabe I Associate Professor

SPEECH RECOGNITION AND UNDERSTANDING IN ADVERSE ENVIRONMENTS.

- Far-Field Speech Recognition
- Multi-Speaker Speech Recognition
- Speaker Diarization
- Speech Enhancement and Separation
- Audio Scene Analysis

DEEP LEARNING FOR AUDIO, SPEECH, AND LANGUAGE PROCESSING

- End-To-End Speech Recognition, Speech Synthesis, And Speech Translation
- End-To-End Integration of Audio, Speech, and Language Processing Modules

Sean Welleck I Assistant Professor

MACHINE LEARNING FOR LANGUAGE TECHOLOGIES

- Large Language Models
- Inference Algorithms
- Learning Algorithms
- Reasoning

AI FOR SCIENCE

- Machine Learning and Mathematics
- Machine Learning and Formal Verification

Chenyan Xiong I Associate Professor

Foundation and Large Language Models
• Efficient Pretraining through data, architecture, and Infrastructure
• Understanding the source of intelligence of foundation models
Next-Gen Information Retrieval
• New ways to connect user and information GenAI.
• Search, RAG, Recommendation, Video Generation, etc.
Bring Revolution of GenAI to Various Verticals
• Healthcare Foundation Models
• Other Interesting Verticals

Yiming Yang I Professor

LARGE LANGUAGE MODELS (LLMs)
• Alignment with Human-valued Principles and Scalable Oversight beyond Human Supervision
MULTIMODAL FOUNDATION MODELS
• LLM-based joint learning & reasoning with text, images, graphs, codes, human feedback, etc.
GENERATIVE AI for SCIENCE
• Generative neural networks with graph-based learning for combinatorial optimization, mathematical reasoning, physics simulation, and more