Language Technologies Institute

Faculty Research Areas

Carnegie Mellon University
School of Computer Science

KEYWORDS AND AREAS OF RESEARCH

NATURAL LANGUAGE PROCESSING

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

COMPUTATIONAL LINGUISTICS

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

SPEECH AND AUDIO PROCESSING

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

MULTIMODAL LEARNING

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

INFORMATION RETRIEVAL

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

MACHINE LEARNING FOR LANGUAGE TECHNOLOGIES

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

KNOWLEDGE-BASED AI

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

FOUNDATION MODEL CORE

- Data-Centric Methodologies (Xiong)
- Efficient Pretraining (Xiong)
- Inference Algorithms (Welleck)
- Post-Training (Mireshghallah, Welleck, Xiong)
- Scaling Up (Xiong)
- Test-time Scaling (Welleck)

AI FOR SCIENCE AND MEDICINE

- AI and Formal Methods (Welleck)
- AI for Drug Discovery (Li)
- AI for Mathematics (Welleck)
- Healthcare AI (Xiong)
- Molecule generation, tool use, and latent modeling (Mireshghallah)
- Protein Modeling and Design (Li)

Core Faculty

Akari Asai I Assistant Professor

https://akariasai.github.io/

LARGE LANGUAGE MODELS AND AGENTS
- Learning and inference-time algorithms for LLMs
- Retrieval and retrieval-time algorithms for LLMs
- Tool-using and agentic systems

RELIABILITY, EVALUATION, AND ALIGNMENT IN LLMS
- Hallucination mitigation, citation fidelity, and copyright risks
- Evaluation frameworks for open-ended, expert-domain tasks
- Training methods for trustworthy behavior and verifiable outputs

AI FOR EXPERT DOMAINS AND UNDERSERVED COMMUNITIES
- Scientific agents for accelerating discovery and decision-making
- Code generation and natural language interfaces to software
- Multilingual and multicultural NLP for low-resource settings

Yonatan Bisk I Assistant Professor

https://talkingtorobots.com/yonatanbisk.html

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?

Ralph Brown I Principal Systems Scientist

https://www.cs.cmu.edu/~ralf/

INFORMATION EXTRACTION
- Text normalization

LANGUAGE IDENTIFICATION
- Identifying short texts in 1000+ languages

DIGITAL FORENSICS
- Text extraction
- Reconstructing corrupted file

Carlos Busso I Professor

https://www.lti.cs.cmu.edu/people/faculty/busso-carlos.html

SPEECH AND MULTIMODAL SYSTEMS
- Affective Computing
- Speech emotion recognition
- Multimodal Machine Learning
- Speech processing
- Social Signal Processing
- Nonverbal behaviors for conversational agents
- Multimodal applications
- Audio Language Models

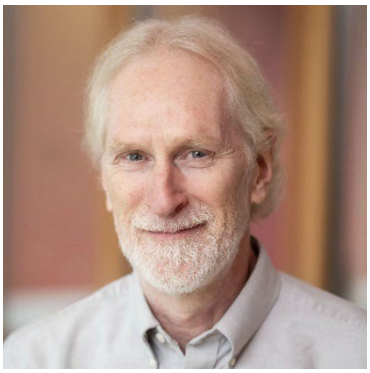

Jamie Callan I Professor & Director, PhD Program

https://www.cs.cmu.edu/~callan/

SEARCH ENGINE ARCHITECTURES
- Discrete and continuous text representations
- Lexical and neural search architectures

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

Mona Diab I Professor I LTI Director I ACL Fellow

https://www.lti.cs.cmu.edu/people/faculty/diab-mona.html

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

https://841.io/

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

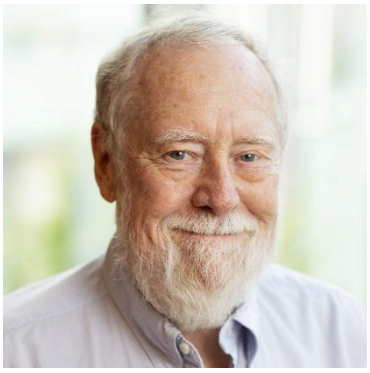

Scott Fahlman I Professor Emeritus, LTI & CSD

https://www.cs.cmu.edu/~sef/

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 to 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 | Assoc. Dean, Doctoral Programs, SCS & Chair, MLT Program

https://www.cs.cmu.edu/~ref/

Daniel Fried | Assistant Professor

https://dpfried.github.io/

GROUNDED INTERACTION

- Grounding language to perception and action
- Language interfaces
- LLM Agents

CODE GENERATION

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

APPLIED PRAGMATICS

- Understanding implicit language
- Contextual language understanding & generation

Alex Hauptmann | Research Professor

https://www.lti.cs.cmu.edu/people/faculty/hauptmann-alexander.html

MULTIMEDIA ANALYSIS AND RETRIEVAL

- Large scale analysis of internet and surveillance video
- Multimedia for healthcare
- Multimedia for human rights

Daphne Ippolito | Assistant Professor

https://www.daphnei.com/

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
- AI for the digital humanities

Alon Lavie I Distinguished Career Professor/Associate Director of MSAII

https://www.cs.cmu.edu/~alavie/

MULTILINGUAL NLP AND LANGUAGE TECHNOLOGIES
- Machine Translation
- Multilingual Content Generation
- Dialogue and Speech Translation
- Translator Assistive Technologies

MULTILINGUAL QUALITY EVALUTAION AND ASSESSMENT
- Automated Translation Quality Evaluation
- Human Translation Quality Evaluation Methods
- Automated Multilingual Chatbot and Dialogue Quality Evaluation

Lori Levin I Research Professor

http://www.cs.cmu.edu/~lsl/

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

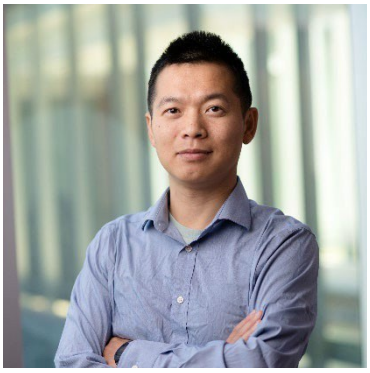

Lei Li I Associate Professor

https://www.cs.cmu.edu/~leili/

LARGE LANGUAGE MODELS
- Security, Safety, Privacy, Copyright of LLM
- LLM Agent, reasoning, code generation
- Efficient LLM systems
- LLM Alignment

MULTILINGUAL NLP
- Machine Translation (how to translate 1000 languages?)
- Speech Translation (how to translate speech to speech in real time?)

AI FOR SCIENCE
- AI for Drug Discovery, design effective small molecule drugs for disease target
- Protein Design
- Agentic Scientific Workflow

Niloofar Mireshghallah I Assistant Professor, LTI, EPP and CyLab

https://mireshghallah.github.io/

- Privacy and security, scientific reasoning

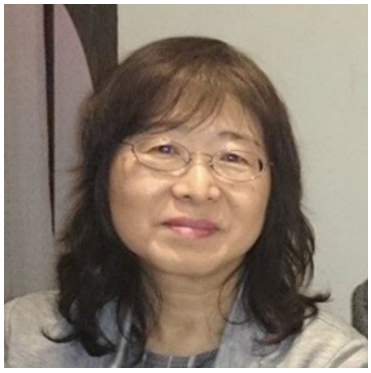

Teruko Mitamura | Research Professor & Director, MIIS Program

https://www.lti.cs.cmu.edu/people/faculty/mitamura-teruko.html

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

MUTLIMODAL AI
- Reactive and Proactive Agentic AI
- Multimodal Procedural Activity Understanding

Louis-Philippe Morency | Associate Professor

https://www.cs.cmu.edu/~morency/

MULTIMODAL MACHINE LEARNING (MULTIMODAL AI)
- Understand, interpret, and process information from multiple heterogeneous, contingent and asynchronous modalities

ARTIFICIAL SOCIAL INTELLIGENCE (SOCIAL AI)
- Analyze, recognize and predict subtle human communicative behaviors during social interactions

HEALTH BEHAVIOR INFORMATICS (HEALTH AI)
- Technologies to support clinical practice during diagnosis and treatment of mental health disorders

David R. Mortensen | Assistant Research Professor

https://www.cs.cmu.edu/~dmortens/

MODELING LANGUAGE AS A DYNAMIC SYSTEM
- Computational models of language variation and change
- Models of language variation in social and geographic space
- Speech and language technologies for people with communicative disabilities

ANALYZING PHONETICS, PHONOLOGY, AND MORPHOLOGY
- Interpreting the phonological and morphological behavior of neural models
- Leveraging phonetic, phonological, and morphological structure for speech and language applications

Graham Neubig | Associate Professor

https://www.phontron.com/

MACHINE LEARING 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 I Professor/Associate Director for Education & Curriculum

https://www.cs.cmu.edu/~ehn/

AUTOMATIC QUESTION ANSWERING

- Multi-strategy approaches for practical real-world applications

CURRICULUM LEARING FOR MULTIMODAL UNDERSTANDING

- Curation of datasets and training algorithms to improve understanding and reasoning in artificial agents

ADVANCED INFORMATION RETRIEVAL

- Improving the performance of sparse and dense retrieval models for practical applications

Kemal Oflazer I Teaching Professor

https://sites.google.com/andrew.cmu.edu/kemal-oflazer/

TEXT SIMPLIFICATION

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

NLP APPLICATIONS

- Generating massive cross-lingual NLI data from parallel corpora

COMPUTATIONAL MORPHOLOGY

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

Adam Paulisick I Distinguished Service Professor

https://www.linkedin.com/in/paulisick/

- Multi-Agent Conversational Architectures for Real-World Decision Making
- Network-Aware Conversational Agents for Distributed Service Discovery
- Contextual Advertising and Agent-Mediated Word-of-Mouth Monetization
- Conversational Competitive Intelligence and Market Sensemaking
- Human–AI Cofounding in AI-Native Ventures

Bhiksha Ramakrishnan I Professor

https://cmu-mlsp.github.io/team/bhiksha_raj

AUDIO ANALYSIS

- Audio content analysis, with applications to acoustic intelligence, surveillance, content-based retrieval
- Speech, Audio and Multimodal foundation models
- 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

https://www.cs.cmu.edu/~cprose/

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

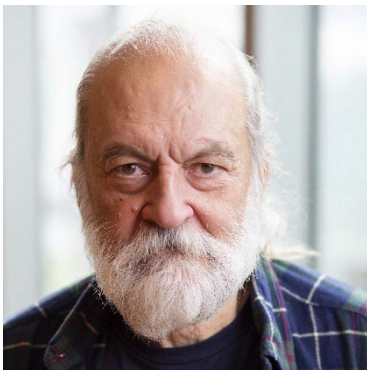

Alexander Rudnicky I Research Professor Emeritus

http://www.cs.cmu.edu/~air/

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

Maarten Sap I Assistant Professor

https://maartensap.com/

SOCIAL INTELLIGENCE IN AI AGENTS
- Assessing Theory-of-Mind, and pragmatics of AI systems
- Unlocking social intelligence and interaction abilities via simulation

MEASURING THE ETHICS AND HARMS OF AI
- Agentic AI safety in multi-turn, tool-using, and user-facing settings
- Human-AI reliance and anthropomorphism
- Alignment biases in AI (culture, race, gender)

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

Michael Shamos I Distinguished Career Professor & Director, MSAII Program

http://euro.ecom.cmu.edu/shamos.html

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

http://ayesha.lti.cs.cmu.edu/mlsp/people/rsingh/index.html

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

Emma Strubell I Assistant Professor/Education & Curriculum Chair

https://strubell.github.io/

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

https://www.cs.cmu.edu/~ahw/

SPEECH TRANSLATION / LANGUAGE TRANSPARENCE

- Lecture Translator
- Speech Translation
- Machine Translators and User Interfaces
- Face Dubbing

NEURAL SPEECH & LANGUAGE

- Neural Robust Speech & Speaker Reco, Translation, Summarization

MUTLIMODAL SYSTEMS

- Emotion, Style, Gestures, Synthesis

Shinji Watanabe I Associate Professor

https://sites.google.com/view/shinjiwatanabe

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 | Assistant Professor

https://wellecks.com/

MACHINE LEARNING FOR LANGUAGE TECHNOLOGIES

- Large Language Models
- Foundations of Generative Models
- AI Agents and Reasoning

AI AND FORMAL METHODS

- AI for Mathematics
- Code Generation

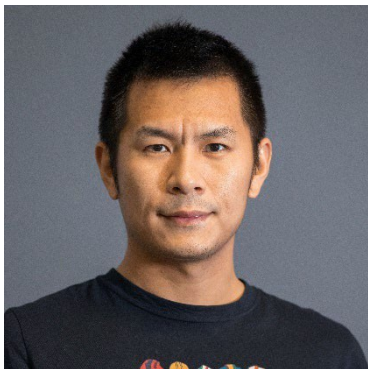

Chenyan Xiong | Associate Professor

https://www.cs.cmu.edu/~cx/

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 | Professor

http://www.cs.cmu.edu/%7Eyiming/

LIKELIHOOD-BASED GENERATIVE REASONING

- Latent reasoning as probabilistic inference, unifying diffusion/flow models and likelihood-based preference optimization for principled alignment

AGENTIC LLMs, RETRIEVAL & EVALUATION

- Community-driven agents, generative retrieval, tool-use RL, scalable test-time reasoning, and rigorous benchmarking

GENERATIVE AI FOR SCIENCE & OPTIMIZATION

- Unsupervised generative methods for combinatorial optimization, MILP, PDE solvers, inverse design, and scientific discovery

Affiliated and Adjunct Faculty

Jeff Bigham I Associate Professor

http://www.cs.cmu.edu/~jbigham/

HUMAN COMPUTER INTERACTION

- Accessibility for People with Disabilities
- Design, Safety & Responsible AI
- Computer Use Agent Collaboration

Roni Rosenfeld I Professor

https://www.ml.cmu.edu/people/bio/roni-rosenfeld/

MACHINE LEARNING APPLICATIONS

- Epidemic detection, tracking and forecasting (Delphi Group)
- AI and Society

Norman Sadeh I Professor, Co-Director Privacy Engineering & AI Governance Program

www.normsadeh.org

Richard Stern I Professor

https://www.lti.cs.cmu.edu/people/faculty/stern-richard.html

SPEECH

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

Sherry Wu I Assistant Professor

https://www.cs.cmu.edu/~sherryw/

Real-World AI Evaluation

- Human-AI collaboration trajectory mining
- AI contribution tracking

Human-centered AIs

- Nuanced instruction following
- Inference time user understanding

Future-of-Work / AI-centered Humans

- AI literacy
- Skill shifts in AI-infused workspace
- Digital twins