

YING SHEN

shenyong19941216@163.com

yingshen-ys.github.io.

EDUCATION

Carnegie Mellon University, School of Computer Science
M.S. in Intelligent Information Systems, GPA: 3.8/4.0

Pittsburgh, PA
Sept 2017 – Dec 2018

Fudan University, School of Software

B.E. in Software Engineering, GPA: 3.68/4.0, Rank: 3/86, National Scholarship (top 1%)

Shanghai, China
Sept 2013 - July 2017

PUBLICATIONS

[1] **Efficient Low-rank Multimodal Fusion with Modality-Specific Factors**

Zhun Liu*, **Ying Shen*** (* indicates equal contribution.), Varun Lakshminarasimhan, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. The 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018), Oral (12.8% acceptance rate).

[2] **Dynamically Adjusting Word Representations Using Nonverbal Behaviours**

Yansen Wang, **Ying Shen**, Zhun Liu, Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019).

WORK EXPERIENCE

Research Associate, Carnegie Mellon University

Pittsburgh, PA

• Advisor: Prof. Louis-Philippe Morency and Prof. Graham Neubig

Jan 2019 – Present

– Researching in deep learning and its applications, especially on representation learning and dialogue understanding.

Graduate Research Assistant, Carnegie Mellon University

Pittsburgh, PA

• Advisor: Prof. Louis-Philippe Morency

Sept 2017 – Jan 2019

– Researching in deep learning and its applications, especially on representation learning and natural language processing.

Research Intern, Carnegie Mellon University

Pittsburgh, PA

• Advisor: Prof. Yoichi Matsuyama and Prof. Justine Cassell

July 2016 - Sept 2016

– Building natural language understanding/generation module for socially aware robot assistant.

RESEARCH EXPERIENCE

• **Learning contextualized utterance representations for dialogue**

Sept 2019 - Present

– Proposed novel dialogue representation model that leverages the speaker identification for learning the contextualized utterance representations of dialogue through unsupervised pre-training methods

• **Refining Absolute Facial Expression Intensities using Within-Person Comparisons**

Jan 2019 – Sept 2019

– Proposed a global-local temporal ranking model in Bayesian framework to derive a person-specific ranking of facial expression intensities based on pairwise comparisons

Unsupervised Learning of Multimodal Representations

Feb 2018 – Dec 2018

– Proposed an unsupervised learning framework to learn multimodal representations that capture not only the rich information on syntactic and semantic relationships between words but also the relation between nonverbal elements and the context in which they occur

– Extended the training objective of the Skip-Gram model from finding word representations that are useful for predicting surrounding words to also capturing the distributions of features extracted from visual and acoustic contexts

– Applied pre-trained multimodal embeddings to other multimodal tasks and showed that multimodal-enriched word embeddings help in achieving better performance

- **Dynamically Adjusting Word Representations Using Nonverbal Behaviors (AAAI-19)** May 2018 - Sept 2018
 - Proposed a neural framework to model the meaning of words and sentences in different nonverbal contexts (e.g. visual and acoustic signals) by considering the nonverbal aspects of human communication as a shift of its verbal representation
 - Achieved competitive performance on two benchmark datasets for multimodal sentiment analysis and emotion recognition; visualizations of the shifted word representations allowed team to gain a better understanding of the impact of subword modeling and dynamic shifts on modeling word meaning
- **Efficient Low-rank Multimodal Fusion with Modality-Specific Factors (ACL-18 Oral)** Oct 2017 - May 2018
 - Proposed an efficient fusion method to integrate multiple unimodal representations (e.g. verbal, visual and audio) into one compact multimodal representation and reduce the computational and model complexity typically present in previous tensor-based fusion approach
 - Designed a low-rank multimodal fusion method to decompose the weight tensor into a set of low-rank modality-specific factors, enabling it to scale linearly as the number of modalities increases
 - Achieved competitive results on three different multimodal tasks: sentiment analysis, speaker trait analysis, and emotion recognition, while drastically reducing computational complexity compared to other methods that utilize tensor representation

PROJECT EXPERIENCE

- **Improving Machine Translation Quality by Cross-lingual Natural Language Inference** Nov 2018 – Dec 2018
 - Proposed a new framework for end-to-end neural machine translation (NMT) by using Cross-lingual Natural Language Inference, a task that is highly dependent on semantics, as a source of error signal in training NMT systems.
 - Applied the minimum risk training method to optimize our NMT system towards generating translations that scores higher in both machine translation and natural language inference tasks.
- **Low-resource Machine Translation** Oct 2018
 - Implemented a Neural Machine Translation model for low-resource translation with Universal Lexical Representation module that representing all tokens of the low-resource source language as a probabilistic mixture of embeddings of tokens from a set of pre-defined “universal tokens”
 - Utilized back-translation to created additional pseudo-parallel data by translating monolingual data with a trained phrase-based statistical machine translation using Moses
- **Dependency Parsing with Deep Reinforcement Learning** Feb 2018 - May 2018
 - Trained a reinforcement learning agent using the Advantage Actor Critic (A2C) algorithm to perform non-greedy decoding with transition-based dependency parser by considering the future rewards
 - Performed experiments on the English Penn Treebank (PTB) datasets and achieved 90.63% unlabeled accuracy, which improved around 0.4% accuracy compared to the supervised neural dependency parser

TEACHING EXPERIENCE

Teaching Assistant, Carnegie Mellon University Fall 2018, 2019
 Course: 11-777, Advanced Multimodal Machine Learning | Instructor: Prof. Louis-Phillipe Morency

AWARDS AND HONORS

KLA-Tencor Excellent Student Scholarship, Fudan University (top 1%)	Oct 2016
EMC Excellent Student Scholarship, Computer Science Department, Fudan University (top 2%)	Apr 2016
National Scholarship (top 1%)	Oct 2015
iShamrock Software Competition 2015, 1st Runner-up	Mar 2015
Excellent Student in Computer Science School, Fudan University (top 1%)	Nov 2014
Fudan University Excellent Student Scholarship, First Prize (top 4%)	Oct 2014

SKILLS AND INTERESTS

Computer Languages	Java, Python, C++/C, MATLAB, SQL, HTML/CSS, JavaScript
Frameworks and Tools	PyTorch, TensorFlow, Keras, Spring, Hibernate, WebGL, OpenCV, AWS, Hadoop, HBase
Languages	Mandarin Chinese (Native), English (Fluent), Japanese (Basic Conversational Proficiency)
Interests	Painting, Swimming, and Traveling