Kai-Wei Chang

University of California, Los Angeles
ENG VI 374
Los Angeles, CA 90095

kw@kwchang.net
http://kwchang.net

RESEARCH INTERESTS

- *Computational approaches to natural language processing.*
- *Tractable machine learning methods for complex and big data.*
- *Trustworthy NLP: Fairness, Accountability, Transparency, and Privacy in AI.*
- *Multimodal representation models.*

EDUCATION AND EXPERIENCE

University of California Los Angeles, CA

Associate Professor, Computer Science	2022 –
Assistant Professor, Computer Science	2017 – 2022

Amazon.com, Inc.

Amazon Visiting Academics	2020 –

University of Virginia, VA

Assistant Professor, Computer Science	2016 – 2017

Microsoft, MA

Postdoctoral researcher, Microsoft Research New England Lab	2015 – 2016

University of Illinois at Urbana-Champaign, IL

Ph.D. in Computer Science	2010 – 2015

Cognitive Computation Group; w/ Dan Roth

National Taiwan University, Taipei, Taiwan

M.S. in Computer Science and Information Engineering	2009

Machine Learning and Data Mining Group; w/ Chih-Jen Lin

B.S. in Computer Science and Information Engineering	2007
B.S. in Electrical Engineering (Dual Degree)	2007

SELECTED AWARDS

Sloan Research Fellowship	2021

The Sloan Research Fellowships seek to stimulate fundamental research by early-career scientists and scholars of outstanding promise.

Google Research Scholar Award	2021
Best Long Paper Award, EMNLP	2017
Best Paper Award, SIGKDD	2010
Best Paper Finalist, CVPR	2022
Okawa Research Grant Award	2018
NSF CRII Award	2016

PUBLICATIONS

26,000+ Google Scholar citations in total. H-index: 51, i10-index: 127.
Published at ACL, EMNLP, NAACL, AAAI, ICML, NeurIPS, KDD, etc.

Refereed Conference Publications

[1] Kuan-Hao Huang, Varun Iyer, I.-Hung Hsu, Anoop Kumar, **Kai-Wei Chang**, and Aram Galstyan. "ParaAMR: A Large-Scale Syntactically Diverse Paraphrase Dataset by AMR Back-Translation." In the Annual Meeting of the Association for Computational Linguistics (ACL 2023).

[2] Liunian Harold Li, Jack Hessel, Youngjae Yu, Xiang Ren, **Kai-Wei Chang**, and Yejin Choi. "Symbolic Chain-of-Thought Distillation: Small Models Can Also 'Think' Step-by-Step." In the Annual Meeting of the Association for Computational Linguistics (ACL 2023).

[3] Tanmay Parekh, I.-Hung Hsu, Kuan-Hao Huang, **Kai-Wei Chang**, and Nanyun Peng. "GENEVA: Pushing the Limit of Generalizability for Event Argument Extraction with 100+ Event Types." In the Annual Meeting of the Association for Computational Linguistics (ACL 2023).

[4] Chenghao Yang, Fan Yin, He He, **Kai-Wei Chang**, Xiaofei Ma, and Bing Xiang. "Efficient Shapley Values Estimation by Amortization for Text Classification." In the Annual Meeting of the Association for Computational Linguistics (ACL 2023).

[5] I.-Hung Hsu, Kuan-Hao Huang, Shuning Zhang, Wenxin Cheng, Prem Natarajan, **Kai-Wei Chang**, and Nanyun Peng. "TAGPRIME: A Unified Framework for Relational Structure Extraction." In the Annual Meeting of the Association for Computational Linguistics (ACL 2023).

[6] Ninareh Mehrabi, Palash Goyal, Apurv Verma, Jwala Dhamala, Varun Kumar, Qian Hu, **Kai-Wei Chang**, Richard Zemel, Aram Galstyan, and Rahul Gupta. "Resolving Ambiguities in Text-to-Image Generative Models." In the Annual Meeting of the Association for Computational Linguistics (ACL 2023).

[7] Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck, and **Kai-Wei Chang**. "A Survey of Deep Learning for Mathematical Reasoning." In the Annual Meeting of the Association for Computational Linguistics (ACL 2023).

[8] Masoud Monajatipoor, Liunian Harold Li, Mozhdeh Rouhsedaghat, Lin Yang, and **Kai-Wei Chang**. "Transferring In-Context Learning Ability From Language Models to Vision-Language Models." In the Annual Meeting of the Association for Computational Linguistics (ACL short), 2023.

[9] Jianfeng Chi, Wasi Uddin Ahmad, Yuan Tian, and **Kai-Wei Chang**. "PLUE: Language Understanding Evaluation Benchmark for Privacy Policies in English." In the Annual Meeting of the Association for Computational Linguistics (ACL short), 2023.

[10] Nikil Roashan Selvam, Sunipa Dev, Daniel Khashabi, Tushar Khot, and **Kai-Wei Chang**. "The Tail Wagging the Dog: Dataset Construction Biases of Social Bias Benchmarks." In the Annual Meeting of the Association for Computational Linguistics (ACL short), 2023.

[11] Yixin Wan, Kuan-Hao Huang, and **Kai-Wei Chang**. "PIP: Parse-Instructed Prefix for Syntactically Controlled Paraphrase Generation." In the Annual Meeting of the Association for Computational Linguistics (ACL short), 2023.

[12] Zixuan Ling, Xiaoqing Zheng, Jianhan Xu, Jinshu Lin, **Kai-Wei Chang**, Cho-Jui Hsieh, and Xuanjing Huang. "Enhancing Unsupervised Semantic Parsing with Distributed Contextual Representations." In *the Annual Meeting of the Association for Computational Linguistics (ACL-Finding 2023)*.

[13] Rui Sun, Zhecan Wang, Haoxuan You, Noel Codella, **Kai-Wei Chang**, and Shih-Fu Chang. "UniFine: A Unified and Fine-grained Approach for Zero-shot Vision-Language Understanding."

In *the Annual Meeting of the Association for Computational Linguistics (ACL-Finding 2023)*.

[14] Wasi Ahmad, Md Golam Rahman Tushar, Saikat Chakraborty, and **Kai-Wei Chang**. "AVATAR: A Parallel Corpus for Java-Python Program Translation." In *the Annual Meeting of the Association for Computational Linguistics (ACL-Finding, short, 2023)*.

[15] Ziniu Hu, Ahmet Iscen, Chen Sun, Zirui Wang, **Kai-Wei Chang**, Yizhou Sun, Cordelia Schmid, David A. Ross, and Alireza Fathi. "REVEAL: Retrieval-Augmented Visual-Language Pre-Training with Multi-Source Multimodal Knowledge." In *the Computer Vision and Pattern Recognition Conference (CVPR 2023)*.

[16] Da Yin, Feng Gao, Govind Thattai, Michael Johnston, and **Kai-Wei Chang**. "GIVL: On Improving Geographical Inclusivity of Vision-and-Language Models with Pre-Training Methods." In *the Computer Vision and Pattern Recognition Conference (CVPR 2023)*.

[17] Honghua Zhang, Liunian Harold Li, Tao Meng, **Kai-Wei Chang**, and Guy Van den Broeck. "On the Paradox of Learning to Reason from Data." In *the International Joint Conferences on Artificial Intelligence (IJCAI 2023)*.

[18] Kareem Ahmed, **Kai-Wei Chang**, and Guy Van den Broeck. "Semantic Strengthening of Neuro-Symbolic Learning." *In the International Conference on Artificial Intelligence and Statistics (AISTATS 2023)*.

[19] Anaelia Ovalle, Arjun Subramonian, Vagrant Gautam, Gilbert Gee, and **Kai-Wei Chang**. "Factoring the Matrix of Domination: A Critical Review and Reimagination of Intersectionality in AI Fairness." In *AAAI/ACM Conference on AI, Ethics, and Society (AIES 2023)*.

[20] Pan Lu, Liang Qiu, **Kai-Wei Chang**, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic Prompt Learning via Policy Gradient for Semi-structured Mathematical Reasoning. *International Conference on Learning Representation (ICLR 2023)*.

[21] Fan Yin, Yao Li, Cho-Jui Hsieh, and **Kai-Wei Chang**. ADDMU: Detection of Far-Boundary Adversarial Examples with Data and Model Uncertainty Estimation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*.

[22] Zhecan Wang, Haoxuan You, Yicheng He, Wenhao Li,**Kai-Wei Chang**, and Shih-Fu Chang. Understanding ME? Multimodal Evaluation for Fine-grained Visual Commonsense In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*.

[23] Ziniu Hu, Yichong Xu, Wenhao Yu, Shuohang Wang, Ziyi Yang, Chenguang Zhu, **Kai-Wei Chang**, and Yizhou Sun. Empowering Language Models with Knowledge Graph Reasoning for Open-Domain Question Answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*.

[24] Da Yin, Hritik Bansal, Masoud Monajatipoor, Liunian Harold Li, and **Kai-Wei Chang**. GeoMLAMA: Geo-Diverse Commonsense Probing on Multilingual Pre-Trained Language Models. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*.

[25] Hritik Bansal, Da Yin, Masoud Monajatipoor, and **Kai-Wei Chang**. How well can Text-to-Image Generative Models understand Ethical Natural Language Interventions?

[26] Ziniu Hu, Yichong Xu, Wenhao Yu, Shuohang Wang, Ziyi Yang, Chenguang Zhu, **Kai-Wei Chang**, and Yizhou Sun. Empowering Language Models with Knowledge Graph Reasoning for Open-Domain Question Answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022, short)*.

[27] Haoxuan You, Rui Sun, Zhecan Wang, **Kai-Wei Chang**, and Shih-Fu Chang. Find Someone Who: Visual Commonsense Understanding in Human-Centric Grounding. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022, Finding)*.

[28] Jieyu Zhao, Xuezhi Wang, Yao Qin, Jilin Chen, and **Kai-Wei Chang**. Investigating Ensemble Methods for Model Robustness Improvement of Text Classifiers. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022, Finding short)*.

[29] Jianfeng Chi, William Shand, Yaodong Yu, **Kai-Wei Chang**, Han Zhao, and Yuan Tian. Conditional Supervised Contrastive Learning for Fair Text Classification. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022, finding)*.

[30] Di Wu, Wasi Uddin Ahmad, Sunipa Dev, and **Kai-Wei Chang**. Representation Learning for Resource-Constrained Keyphrase Generation In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2022, finding)*.

[31] Jwala Dhamala, Varun Kumar, Rahul Gupta, **Kai-Wei Chang**, Aram Galstyan. An analysis of the effects of decoding algorithms on fairness in open-ended language generation. In IEEE Spoken Language Technology Workshop (SLT 2022)

[32] Arjun Subramonian, **Kai-Wei Chang**, and Yizhou Sun. On the Discrimination Risk of Mean Aggregation Feature Imputation in Graphs In *Neural Information Processing Systems (NeurIPS 2022)*.

[33] Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, **Kai-Wei Chang**, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to Explain: Multimodal Reasoning via Thought Chains for Science Question Answering. In *Neural Information Processing Systems (NeurIPS 2022)*.

[34] Tao Meng, Sidi Lu, Nanyun Peng, and **Kai-Wei Chang**. Controllable Text Generation with Neurally-Decomposed Oracle In *Neural Information Processing Systems (NeurIPS 2022)*. **Oral**

[35] Kareem Ahmed, Stefano Teso, **Kai-Wei Chang**, Guy Van den Broeck, and Antonio Vergari. Semantic Probabilistic Layers for Neuro-Symbolic Learning. In *Neural Information Processing Systems (NeurIPS 2022)*.

[36] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, **Kai-Wei Chang**, and Jianfeng Gao, Grounded Language-Image Pre-training, *the Computer Vision and Pattern Recognition Conference (CVPR 2022)*. **Best Paper Finalist**

[37] Kareem Ahmed, Eric Wang, **Kai-Wei Chang**, and Guy Van den Broeck, Neuro-Symbolic Entropy Regularization. *Uncertainty in Artificial Intelligence (UAI 2022)*.

[38] Vijit Malik, Sunipa Dev, Akihiro Nishi, Nanyun Peng, and **Kai-Wei Chang**, Socially Aware Bias Measurements for Hindi Language Representations. *North American Chapter of the Association for Computational Linguistics (NAACL short 2022)*.

[39] I.-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, **Kai-Wei Chang**, and Nanyun Peng. DEGREE: A Data-Efficient Generative Event Extraction Model. *North American Chapter of the Association for Computational Linguistics (NAACL 2022)*.

[40] Kuan-Hao Huang, I.-Hung Hsu, Prem Natarajan, **Kai-Wei Chang**, and Nanyun Peng, Multilingual Generative Language Models for Zero-Shot Cross-Lingual Event Argument Extraction *the Annual Meeting of the Association for Computational Linguistics (ACL 2022)*

[41] Satyapriya Krishna, Rahul Gupta, Apurv Verma, Jwala Dhamala, Yada Pruksachatkun, and **Kai-Wei Chang**, Measuring Fairness of Text Classifiers via Prediction Sensitivity. *the Annual Meeting of the Association for Computational Linguistics (ACL 2022)*

[42] Fan Yin, Zhouxing Shi, Cho-Jui Hsieh, and **Kai-Wei Chang**, On the Sensitivity and Stability of Model Interpretations. *the Annual Meeting of the Association for Computational Linguistics (ACL 2022)*

[43] Yang Trista Cao, Yada Pruksachatkun, **Kai-Wei Chang**, Rahul Gupta, Varun Kumar, Jwala Dhamala, and Aram Galstyan, On the Intrinsic and Extrinsic Fairness Evaluation Metrics for Contextualized Language Representations. *the Annual Meeting of the Association for Computational Linguistics (ACL 2022, short)*

[44] Jianhan Xu, Cenyuan Zhang, Xiaoqing Zheng, Linyang Li, Cho-Jui Hsieh, **Kai-Wei Chang**, and Xuanjing Huang. Towards Adversarially Robust Text Classifiers by Learning to Reweight Clean Examples. *the Annual Meeting of the Association for Computational Linguistics (ACL-Finding 2022)*

[45] Umang Gupta, Jwala Dhamala, Varun Kumar, Apurv Verma, Yada Pruksachatkun, Satyapriya Krishna, Rahul Gupta, **Kai-Wei Chang**, Greg Ver Steeg, and Aram Galstyan, Mitigating Gender Bias in Distilled Language Models via Counterfactual Role Reversal. *the Annual Meeting of the Association for Computational Linguistics (ACL-Finding 2021)*

[46] Cenyuan Zhang, Xiang Zhou, Yixin Wan, Xiaoqing Zheng, **Kai-Wei Chang**, and Cho-Jui Hsieh, Improving the Adversarial Robustness of NLP Models by Information Bottleneck. *the Annual Meeting of the Association for Computational Linguistics (ACL-Finding 2021)*

[47] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, **Kai-Wei Chang**, Zhewei Yao, and Kurt Keutz, How Much Can CLIP Benefit Vision-and-Language Tasks? *International Conference on Learning Representation (ICLR 2022).*

[48] Zhecan Wang, Haoxuan You, Liunian Harold Li, Alireza Zareian, Suji Park, Yiqing Liang, **Kai-Wei Chang**, and Shih-Fu Chang, SGEITL: Scene Graph Enhanced Image-Text Learning for Visual Commonsense Reasoning. *AAAI Conference on Artificial Intelligence (AAAI 2022).*

[49] D. Yin, L. Harold Li, Z. Hu, N. Peng, and **K.-W. Chang**, Broaden the Vision: Geo-Diverse Visual Commonsense Reasoning. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2021).*

[50] S. Dev, M. Monajatipoor, A. Ovalle, A. Subramonian, J. Phillips, and **K.-W. Chang**, Harms of Gender Exclusivity and Challenges in Non-Binary Representation in Language Technologies. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2021).*

[51] L. Yuan, X. Zheng, Y. Zhou, C.-J. Hsieh, and **K.-W. Chang**. On the Transferability of Adversarial Attacks against Neural NLP Models *Conference on Empirical Methods in Natural Language Processing (EMNLP 2021).*

[52] Z. Li, J. Xu, J. Zeng, L. Li, X. Zheng, Q. Zhang, **K.-W. Chang**, and C.-J. Hsieh. Searching for an Effiective Defender: Benchmarking Defense against Adversarial Word Substitution. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2021).*

[53] K.-H. Huang, W. Ahmed, N. Peng, and **K.-W. Chang**, Improving Zero-Shot Cross-Lingual Transfer Learning via Robust Training. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2021).*

[54] Z. Hu, Y. Sun, and **K.-W. Chang**, Relation-Guided Pre-Training for Open-Domain Question Answering. *Conference on Empirical Methods in Natural Language Processing (EMNLP finding 2021).*

[55] R. Parvez, W. Ahmad, S. Chakraborty, B. Ray, and **K.-W. Chang**, Retrieval Augmented Code Generation and Summarization. *Conference on Empirical Methods in Natural Language Processing (EMNLP finding 2021).*

[56] T. Meng and **K.-W. Chang**, An Integer Linear Programming Framework for Mining Constraints from Data. *the 32nd International Conference on Machine Learning (ICML 2021).*

[57] W. Ahmad, X. Bai, S. Lee, and **K.-W. Chang**, Select, Extract and Generate: Neural Keyphrase

Generation with Layer-wise Coverage Attention. *the Annual Meeting of the Association for Computational Linguistics (ACL 2021)*

[58] W. Ahmad, J. Chi, T. Le, T. Norton, Y. Tian, **K.-W. Chang**, Intent Classification and Slot Filling for Privacy Policies. *the Annual Meeting of the Association for Computational Linguistics (ACL 2021)*

[59] W. Ahmad, H. Li, **K.-W. Chang**, Yashar Mehdad, Syntax-augmented Multilingual BERT for Cross-lingual Transfer. *the Annual Meeting of the Association for Computational Linguistics (ACL 2021)*

[60] E. Sheng, **K.-W. Chang**, P. Natarajan, N. Peng, Societal Biases in Language Generation: Progress and Challenges *the Annual Meeting of the Association for Computational Linguistics (ACL 2021)*

[61] Y. Zhou, X. Zheng, C.-J. Hsieh, **K.-W. Chang**, X. Huang, Defense against Synonym Substitution-based Adversarial Attacks via Dirichlet Neighborhood Ensemble. *the Annual Meeting of the Association for Computational Linguistics (ACL 2021)*

[62] Y. Pruksachatkun, S. Krishna, J. Dhamala, R. Gupta, **K.-W. Chang**, Does Robustness Improve Fairness? Approaching Fairness with Word Substitution Robustness Methods for Text Classification. *the Annual Meeting of the Association for Computational Linguistics (ACL 2021, Finding)*

[63] J. Zhao, D. Khashabi, T. Khot, A. Sabharwal **K.-W. Chang**. Ethical-Advice Taker: Do Language Models Understand Natural Language Interventions? *the Annual Meeting of the Association for Computational Linguistics (ACL 2021, short, Finding)*

[64] R. Parvez, **K.-W. Chang**. Evaluating the Values of Sources in Transfer Learning. *North American Chapter of the Association for Computational Linguistics (NAACL 2021).*

[65] W. Ahmad, S. Chakraborty, B. Ray, **K.-W. Chang**. Unified Pre-training for Program Understanding and Generation. *North American Chapter of the Association for Computational Linguistics (NAACL 2021).*

[66] E. Sheng, **K.-W. Chang**, P. Natarajan, N. Peng. "Nice Try, Kiddo": Ad Hominems in Dialogue Systems. *North American Chapter of the Association for Computational Linguistics (NAACL 2021).*

[67] C. Zhang, J. Zhao, H. Zhang, **K.-W. Chang**, C.-J. Hsieh. Double Perturbation: On the Robustness of Robustness and Counterfactual Bias Evaluation. *North American Chapter of the Association for Computational Linguistics (NAACL 2021).*

[68] L. Li, H. You, Z. Wang, A. Zareian, S.-F. Chang, **K.-W. Chang**. Unsupervised Vision-and-Language Pre-training Without Parallel Images and Captions *North American Chapter of the Association for Computational Linguistics (NAACL 2021).*

[69] J. Huang, K.-H. Huang, and **K.-W. Chang**. Disentangling Semantics and Syntax in Sentence Embeddings with Pre-trained Language Models. *North American Chapter of the Association for Computational Linguistics (NAACL 2021).*

[70] A. Uppunda, S. Cochran, J. Foster, A. Arseniev-Koehler, V. Mays, **K.-W. Chang**. Adapting Coreference Resolution for Processing Violent Death Narratives. *North American Chapter of the Association for Computational Linguistics (NAACL 2021).*

[71] J. Dhamala, T. Sun, V. Kumar, S. Krishna, Y. Pruksachatkun, **K.-W. Chang**, R. Gupta. BOLD: Dataset and metrics for measuring biases in open-ended language generation. *in FAccT*, 2021.

[72] K.-H. Huang and **K.-W. Chang**. Generating Syntactically Controlled Paraphrases without Using Annotated Parallel Pairs. *European Chapter of the Association for Computational Linguistics (EACL 2021)*

[73] Y. Zhou, Y. Yan, R. Han, H. Caufield, **K.-W. Chang**, Y. Sun, P. Ping, W. Wang. Clinical Temporal Relation Extraction with Probabilistic Soft Logic Regularization and Global Inference. *AAAI Conference on Artificial Intelligence (AAAI 2021)*.

[74] W. Ahmad, N. Peng, **K.-W. Chang**. GATE: Graph Attention Transformer Encoder for Cross-lingual Relation and Event Extraction. *AAAI Conference on Artificial Intelligence (AAAI 2021)*.

[75] K. Xu, Z. Shi, H. Zhang, Y. Wang, **K.-W. Chang**, M. Huang, B. Kailkhura, X. Lin, and C.-J. Hsieh, Provable, Scalable and Automatic Perturbation Analysis on General Computational Graphs. *Neural Information Processing Systems (NeurIPS 2020)*.

[76] K.-H. Huang, C. Li, **K.-W. Chang**. Generating Sports News from Live Commentary: A Chinese Dataset for Sports Game Summarization. *Asia Chapter of the Association for Computational Linguistics (AACL 2020)*

[77] Z. Hu, Y. Dong, K. Wang, **K.-W. Chang**, and Y. Sun. GPT-GNN: Generative Pre-Training of Graph Neural Networks. *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2020)*.

[78] E. Sheng, **K.-W. Chang**, P. Natarajan, N. Peng. Towards Controllable Biases in Language Generation. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2020, finding)*.

[79] L. Liu, Y. Zhou, J. Xu, X. Zheng, **K.-W. Chang**, X. Huang. Cross-Lingual Dependency Parsing by POS-Guided Word Reordering. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2020, finding)*.

[80] J. Zhao and **K.-W. Chang**. LOGAN: Local Group Bias Detection by Clustering. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2020, short)*.

[81] W. Ahmad, J. Chi, Y. Tian and **K.-W. Chang**. PolicyQA: A Reading Comprehension Dataset for Privacy Policies. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2020, short finding)*.

[82] F. Yin, Q. Long, T. Meng, and **K.-W. Chang**, On the Robustness of Language Encoders against Grammatical Errors. *the Annual Meeting of the Association for Computational Linguistics (ACL 2020)*

[83] D. Yin, T. Meng, and **K.-W. Chang**, SentiBERT: An Effective, Transferable and Interpretable Architecture for Compositional Sentiment Semantics. *the Annual Meeting of the Association for Computational Linguistics (ACL 2020)*

[84] J. Zhao, S. Mukherjee, S. Hosseini, **K.-W. Chang**, and A. H. Awadallah, Gender Bias in Multilingual Embeddings and Cross-Lingual Transfer. *the Annual Meeting of the Association for Computational Linguistics (ACL 2020)*

[85] A. Gaut, T. Sun, S. Tang, Y. Huang, J. Qian, M. ElSherief, J. Zhao, D. Mirza, E. Belding, **K.-W. Chang**, and W. Wang, Towards Understanding Gender Bias in Relation Extraction. *the Annual Meeting of the Association for Computational Linguistics (ACL 2020)*

[86] Y. Zhou, J.-Y. Jiang, J. Zhao, **K.-W. Chang**, and W. Wang, "The Boating Store Had Its Best Sail Ever": Pronunciation-attentive Contextualized Pun Recognition. *the Annual Meeting of the Association for Computational Linguistics (ACL 2020)*

[87] Y. S. Jia, T. Meng, J. Zhao, and **K.-W. Chang**, Mitigating Gender Bias Amplification in Distribution by Posterior Regularization *the Annual Meeting of the Association for Computational Linguistics (ACL 2020, short)*

[88] W. Ahmad, S. Chakraborty, B. Ray, and **K.-W. Chang**, A Transformer-based Approach for Source Code Summarization. *the Annual Meeting of the Association for Computational Linguistics (ACL 2020, short)*

[89] L. Li, M. Yatskar, D. Yin, C.-J. Hsieh, and **K.-W. Chang**, What Does BERT with Vision Look At? *the Annual Meeting of the Association for Computational Linguistics (ACL 2020, short)*

[90] Z. Shi, H. Zhang, **K.-W. Chang**, M. Huang, and C.-J. Hsieh, Robustness Verification for Transformers. *International Conference on Learning Representation (ICLR 2020).*

[91] W. Ahmad, Z. Zhang, X. Ma, **K.-W. Chang**, N. Peng. Cross-lingual Dependency Parsing with Unlabeled Auxiliary Languages. *The SIGNLL Conference on Computational Natural Language Learning (CoNLL 2019).*

[92] M. Chen, Y. Tian, H. Chen, **K.-W. Chang**, S. Skiena, C. Zaniolo. Learning to Represent Bilingual Dictionaries. *The SIGNLL Conference on Computational Natural Language Learning (CoNLL 2019).*

[93] T. Meng, N. Peng, **K.-W. Chang**. Target Language-Aware Constrained Inference for Cross-lingual Dependency Parsing. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2019).*

[94] Y. Zhou, J.-Y. Jiang, **K.-W. Chang**, W. Wang. Learning to Discriminate Perturbations for Blocking Adversarial Attacks in Text Classification. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2019).*

[95] P. Zhou, W. Shi, J. Zhao, K.-H. Huang, M. Chen, R. Cotterell, **K.-W. Chang**. Examining Gender Bias in Languages with Grammatical Gender. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2019).*

[96] R. Parvez, T. Bolukbasi, **K.-W. Chang**, V. Saligrama. Robust Text Classifier on Test-Time Budgets. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2019, short).*

[97] W. Shi, M. Chen, P. Zhou, **K.-W. Chang**. Retrofitting Contextualized Word Embeddings with Paraphrases. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2019, short).*

[98] E. Sheng, **K.-W. Chang**, P. Natarajan, N. Peng. The Woman Worked as a Babysitter: On Biases in Language Generation. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2019, short).*

[99] C. Xia, H. Zhang, J. Moghtader, A. Wu, **K.-W. Chang**. Visualizing Trend of Key Roles in News Articles. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2019, demo).*

[100] T. Wang, J. Zhao, M. Yatskar, **K.-W. Chang**, V. Ordonez. Balanced Datasets Are Not Enough: Estimating and Mitigating Gender Bias in Deep Image Representations. *International Conference on Computer Vision (ICCV 2019).*

[101] W. Ahmad, **K.-W. Chang**, H. Wang. Context Attentive Document Ranking and Query Suggestion. *International ACM SIGIR Conference on Research and Development in Information Retrieval. (SIGIR 2019).*

[102] Z. Hu, T. Chen, **K.-W. Chang**, Y. Sun. Few-Shot Representation Learning for Out-Of-Vocabulary Words. *the Annual Meeting of the Association for Computational Linguistics (ACL 2019).*

[103] T. Sun, A. Gaut, S. Tang, Y. Huang, M. ElSherief, J. Zhao, D. Mirza. **K.-W. Chang**, W. Wang. Debiasing Gender in Natural Language Processing: Literature Review. *the Annual Meeting of the Association for Computational Linguistics (ACL 2019).*

[104] W. Ahmad, Z. Zhang, X. Ma, E. Hovy, **K.-W. Chang**, N. Peng. On Difficulties of Cross-Lingual Transfer with Order Differences: A Case Study on Dependency Parsing. *North American Chapter of the Association for Computational Linguistics (NAACL 2019).*

[105] J. Zhao, T. Wang, M. Yatskar, R. Cotterell, V. Ordonez, **K.-W. Chang**. Gender Bias in Contextualized Word Embeddings. *North American Chapter of the Association for Computational Linguistics (NAACL 2019, short)*.

[106] M. Chen, C. Ju, G. Zhou, X. Chen, T. Zhang, **K.-W. Chang**, C. Zaniolo, W. Wang. Multifaceted Protein-Protein Interaction Prediction Based on Siamese Residual RCNN. *ISMB 2019*.

[107] J. Zhao, Y. Zhou, Z. Li, W. Wang, **K.-W. Chang** Learning Gender-Neutral Word Embeddings *Conference on Empirical Methods in Natural Language Processing (EMNLP 2018, short)*.

[108] M. Alzantot, Y. Sharma, A. Elgohary, B.-J. Ho, M. Srivastava, **K.-W. Chang**. Generating Natural Language Adversarial Examples. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2018, short)*

[109] R. Parvez, S. Chakraborty, B. Ray, **K.-W. Chang**. Building Language Models for Text with Named Entities. *the Annual Meeting of the Association for Computational Linguistics (ACL 2018)*.

[110] W. Ahmad, **K.-W. Chang**, Hongning Wang. Intent-aware Query Obfuscation for Privacy Protection in Personalized Web Search *International ACM SIGIR Conference on Research and Development in Information Retrieval. (SIGIR 2018)*.

[111] C. Jiang, H.-F Yu, C.-J. Hsieh, **K.-W. Chang**, Learning Word Embeddings for Low-resource Languages by PU Learning. *North American Chapter of the Association for Computational Linguistics (NAACL 2018)*.

[112] J. Zhao, T. Wang, M. Yatskar, V. Ordonez, **K.-W. Chang**, Gender Bias in Coreference Resolution:Evaluation and Debiasing Methods. *North American Chapter of the Association for Computational Linguistics (NAACL 2018, short)*.

[113] W. Ahmad, **K.-W. Chang**, A Corpus to Learn Refer-to-as Relations for Nominals. *Language Resources and Evaluation Conference (LREC 2018)*.

[114] M. Chen, Y. Tian, **Kai-Wei Chang**, S. Skiena, C. Zaniolo Co-training Embeddings of Knowledge Graphs and Entity Descriptions for Cross-lingual Entity Alignment *the 22nd International Joint Conferences on Artificial Intelligence (IJCAI 2018)*

[115] S. Preum, R. Parvez, **K.-W. Chang**, J. Stankovic, A Corpus of Drug Usage Guidelines Annotated with Type of Advice. *Language Resources and Evaluation Conference (LREC 2018)*.

[116] L. Feng, M. Ghasemi, **K.-W. Chang**, U. Topcu. Counterexamples for Robotic Planning Explained in Structured Natural Language *International Conference on Robotics and Automation (ICRA 2018)*.

[117] W. Ahmad, **K.-W. Chang**, H. Wang. Multi-Task Learning for Document Ranking and Query Suggestion. *International Confernece on Learning Representation (ICLR 2018)*.

[118] K. Arnold, **K.-W Chang**, A. Kalai, Counterfactual Language Model Adaptation for Suggesting Phrases. *International Joint Conference on Natural Language Processing* (IJCNLP 2017).

[119] J. Zhao, T. Wang, M. Yatskar, V. Ordonez, **K.-W. Chang**. Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*, **Best Long Paper Award**.

[120] T. Bolukbasi, **K.-W. Chang**, Joseph Wang, Venkatesh Saligrama. Structured Prediction with Test-time Budget Constraints. *Thirty-First AAAI Conference on Artificial Intelligence (AAAI 2017)*.

[121] T. Bolukbasi, **K.-W Chang** , James Zou, Venkatesh Saligrama, Adam Kalai, Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. *Neural Information Processing Systems (NeurIPS 2016)*.

[122] K.-W. Chang, H. He, H. Daume III, J. Langford, S. Ross A Credit Assignment Compiler for Joint Prediction. *Neural Information Processing Systems (NeurIPS 2016).*

[123] S. Upadhyay, M. Chang, **K.-W. Chang**, W.-t. Yih, Learning from Explicit and Implicit Supervision Jointly For Algebra Word Problems, *Conference on Empirical Methods in Natural Language Processing (EMNLP 2016).*

[124] **K.-W. Chang**, A. Krishnamurthy, A. Agarwal, H. Daumé III, J. Langford. Learning to search better than your teacher, *the 32nd International Conference on Machine Learning (ICML 2015).*

[125] H. Peng, **K.-W. Chang**, D. Roth. A joint framework for coreference resolution and mention head detection, *The SIGNLL Conference on Computational Natural Language Learning (CoNLL 2015).*

[126] **K.-W. Chang**, S. Upadhyay, G. Kundu and D. Roth Structural learning with amortized inference *The Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015).*

[127] **K.-W. Chang**, W.-t. Yih, B. Yang and C. Meek. Typed tensor decomposition of knowledge bases for relation extraction. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014).*

[128] R. Samdani, **K.-W. Chang**, D. Roth. A discriminative latent variable model for online clustering. *the 31st International Conference on Machine Learning (ICML 2014).* .

[129] **K.-W. Chang**, R. Samdani, D. Roth. A constrained latent variable model for coreference resolution. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2013).*

[130] **K.-W. Chang**, W.-t. Yih, C. Meek. Multi-relational latent semantic analysis. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2013).*

[131] **K.-W. Chang**, V. Srikumar, D. Roth. Multi-core structural SVM training. *European Conference on Machine Learning (ECML 2013).*

[132] **K.-W. Chang**, S. Sundararajan, S. S. Keerthi. Tractable semi-supervised learning of complex structured prediction models. *European Conference on Machine Learning (ECML 2013).*

[133] **K.-W. Chang**, B. Deka, W.-M. H. Hwu, D. Roth. Efficient pattern-based time series classification on GPU. *2012 IEEE 12th International Conference on Data Mining (ICDM 2012).*

[134] **K.-W. Chang** and D. Roth, Selective block minimization for faster convergence of limited memory large-scale linear models. *the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2011).*

[135] H.-F. Yu, C.-J. Hsieh, **K.-W. Chang**, and C.-J. Lin, Large linear classification when data cannot fit in memory, *the 22nd International Joint Conferences on Artificial Intelligence (IJCAI 2011, the Best Paper Track).*

[136] H.-F. Yu, C.-J. Hsieh, **K.-W. Chang**, and C.-J. Lin, Large linear classification when data cannot fit in memory, *the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2010),* **Best research paper**.

[137] F.-L. Huang, C.-J. Hsieh, **K.-W. Chang**, and C.-J. Lin, Iterative scaling and coordinate descent methods for maximum entropy models, *the 47th Annual Meeting of the Association for Computational Linguistics (ACL 2009, short paper).*

[138] S. S. Keerthi, S. Sundararajan, **K.-W. Chang**, C.-J. Hsieh, and C.-J. Lin, A sequential dual method for large scale multi-class linear SVMs, *the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2008).*

[139] C.-J. Hsieh, **K.-W. Chang**, C.-J. Lin, S. Sathiya Keerthi, and S. Sundararajan, A dual coordinate descent method for large-scale linear SVM, *the 25th International Conference on Machine Learning (ICML 2008).*

Refereed Journal Publications

[140] A. Arseniev-Koehler, S. D. Cochran, V. M. Mays, **K.-W. Chang**, and J. G. Foster. Integrating topic modeling and word embedding to characterize violent deaths, *Proceedings of the National Academy of Sciences* 2022.

[141] A. Arseniev-Koehler, J. G. Foster, V. M. Mays, **K.-W. Chang**, and S. D. Cochran. Aggression, Escalation, and Other Latent Themes in Legal Intervention Deaths of Non-Hispanic Black and White Men: Results From the 2003 - 2017 National Violent Death Reporting System, *American Journal of Public Health, 2021.*

[142] C.-P. Li, **K.-W. Chang**, Distributed Block-diagonal Approximation Methods for Regularized Empirical Risk Minimization, *Machine Learning Journal, 2019.*

[143] L. Li, P. Chen, C.-J. Hsieh, **K.-W. Chang**, Efficient Contextual Representation Learning With Continuous Outputs, *Transactions of the Association for Computational Linguistics, 2019*

[144] D. Duong, W. Ahmad, E. Eskin, **K.-W. Chang**, J. Li. Word and sentence embedding tools to measure semantic similarity of Gene Ontology terms by their definitions *Journal of Computational Biology 2018*

[145] H.-F. Yu, C.-J. Hsieh, **K.-W. Chang**, and C.-J. Lin, Large linear classification when data cannot fit in memory, *ACM Transactions on Knowledge Discovery from Data (TKDD) 5(4):23, 2012*

[146] G.-X. Yuan, **K.-W. Chang**, C.-J. Hsieh, C.-J. Lin, A comparison of optimization methods for large-scale L1-regularized linear classification. *Journal of Machine Learning Research 11 (JMLR), 3183-3234, 2010.*

[147] Y.-W. Chang, C.-J. Hsieh, **K.-W. Chang**, Michael Ringgaard, and C.-J. Lin, Training and Testing Low-degree Polynomial Data Mappings via Linear SVM *Journal of Machine Learning Research 11 (JMLR), 1471-1490, 2010.*

[148] F.-L. Huang, C.-J. Hsieh, **K.-W. Chang**, and C.-J. Lin, Iterative scaling and coordinate descent methods for maximum entropy models. *Journal of Machine Learning Research 11 (JMLR), 815-848, 2010.*

[149] R.-E. Fan, **K.-W. Chang**, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research 9 (JMLR), 1871-1874, 2008.*

[150] **K.-W. Chang**, C.-J. Hsieh, and C.-J. Lin, Coordinate descent method for large-scale L2-loss linear SVM. *Journal of Machine Learning Research 9 (JMLR), 1369-1398, 2008.*

Refereed Shared Task System Papers

[151] Kareem Ahmed, Tao Li, Thy Ton, Quan Guo, **Kai-Wei Chang**, Parisa Kordjamshidi, Vivek Srikumar, Guy Van den Broeck, and Sameer Singh.PYLON: A PyTorch Framework for Learning with Constraints. *AAAI Conference on Artificial Intelligence (AAAI demo 2022).*

[152] A. Rozovskaya, **K.-W. Chang**, D. Roth. The Illinois-Columbia System in the CoNLL-2014 Shared Task *Proceedings of the Eighteenth Conference on Computational Natural Language Learning (CoNLL 2014)* (**1st place**).

[153] A. Rozovskaya, **K.-W. Chang**, M. Sammons, D. Roth. The University of Illinois System in the CoNLL-2013 Shared Task. *Proceedings of the Seventeenth Conference on Computational Natural Language Learning (CoNLL) 2013* (**1st place**).

[154] X. Cheng, B. Chen, R. Samdani, **K-W. Chang**, Z. Fei, M. Sammons, J. Wieting, S. Roy, C. Wang, and D. Roth, Illinois Cognitive Computation Group UI-CCG TAC 2013 Entity Linking and Slot Filler Validation Systems, *Text Analysis Conference (TAC 2013).*

[155] **K.-W. Chang**, R. Samdani, A. Rozovskaya, M. Sammons and D. Roth, Illinois-Coref: The UI System in the CoNLL-2012 Shared Task *Proceedings of the Sixteenth Conference on*

Computational Natural Language Learning (CoNLL) 2012 **(4th place in the English closed track).**

[156] **K.-W. Chang**, R. Samdani, A. Rozovskaya, N. Rizzolo, M. Sammons and D. Roth, Inference Protocols for Coreference Resolution. *Proceedings of the Fifteenth Conference on Computational Natural Language Learning (CoNLL) 2011* **(3rd place in the closed track).**

[157] H.-Y. Lo, **K.-W. Chang**, S.-T. Chen, T.-H. Chiang, C.-S. Ferng, C.-J. Hsieh, Y.-K. Ko, T.-T. Kuo, H.-C. Lai, K.-Y. Lin, C.-H. Wang, H.-F. Yu, C.-J. Lin, H.-T. Lin and S.-D. Lin. An ensemble of three classifiers for KDD Cup 2009: expanded linear model, heterogeneous boosting, and selective naive Bayes. *Proceedings of KDD-Cup 2009 competition, vol. 7 of JMLR Workshop and Conference Proceedings, 57-64, 2009.* **(3rd Place out of 400+ submissions in the Slow Track).**

Thesis

[158] K.-W. Chang, Selective algorithms for large-scale classification and structured learning (Ph.D.)

[159] K.-W. Chang, A dual coordinate descent method for large-scale linear SVM (Master)

Workshop Publications and Preprints

[160] Hritik Bansal, Nishad Singhi, Yu Yang, Fan Yin, Aditya Grover, and **Kai-Wei Chang**. "CleanCLIP: Mitigating Data Poisoning Attacks in Multimodal Contrastive Learning." In *ICLR workshop on Reliable and Trustworthy Large Scale Machine Learning, 2023* Best Paper .

[161] Z. Deng, W. Shi, P. Zhou, M. Chen, **K.-W. Chang**. Computational Analysis of French-origin Reborrowing Process for English Loanwords. *ICDM Workshop on Multilingual Cognitive Services*, 2019.

[162] Z. Hu, C. Fan, T. Chen, **K.-W. Chang**, Y. Sun. Pre-Training Graph Neural Networks for Generic Structural Feature Extraction. *ICLR Workshop: Representation Learning on Graphs and Manifolds*, 2019.

[163] W. Shi, M. Chen, Y. Tian, **K.-W. Chang**. Learning Bilingual Word Embeddings Using Lexical Definitions *ACL Representation learning for NLP Workshop 2019.*

[164] S. Upadhyay, **K.-W. Chang**, M. Taddy, A. Kalai, J. Zou. Beyond Bilingual: Multi-senseWord Embeddings using Multilingual Context. *ACL Representation learning for NLP Workshop 2017,* **Best Paper Award**.

[165] K. Arnold, **K.-W Chang**, A Kalai. Learning to Suggest Phrases. *AAAI Workshop on Human-Aware AI Workshop*, 2017.

[166] C.-p. Lee, K.-W. Chang, S. Upadhyay, D. Roth. Distributed Training of Structured SVM. *NeurIPS Workshop on Optimization for Machine Learning*, 2015.

[167] K. -W. Chang, H. Daumé III, J. Langford, S. Ross. Efficient Programmable Learning to Search. *ICML Workshop on Machine Learning System*, 2015.

[168] R. Samdani, **K.-W. Chang**, D. Roth. A Discriminative Latent Variable Model for Clustering of Streaming Data with Application to Coreference Resolution. *ICML workshop on Inferning: Interactions between Inference and Learning*, 2013.

[169] H.-F. Yu, C.-J. Hsieh, **K.-W. Chang**, and C-J. Lin, Pascal Challenge: Linear Support Vector Machines. *Pascal Large Scale Learning Challenge in ICML 2008 Workshop*, 2008.

Patents

[170] Efficient polynomial mapping of data for use with linear support vector machines, Y.-W. Chang, C.-J. Hsieh, **K.-W. Chang**, M. Ringgaard, C.-J. Lin, 2013.

[171] Interactive Context-Based Text Completions, K. Arnold, **K.-W Chang**, A Kalai, 2016 (under review).

FUNDING

I've secured more than $3,000,000 research fund solely the portion for my group from DARPA, NSF, and industrial partners.

PIPP Phase 1: An End-to-End Pandemic Early Warning System by Harnessing Open-Source Intelligence 2022-2024
NSF PIPP. Senior Personnel. PI: Wei Wang

Cisco Research Grant 2023

Optum Research Grant 2023

Cisco Research Grant 2022

Google Research Scholar 2021

Amazon Research Award 2020–2021

AI-DCL: Governing Bias in AI System with Humans in the Decision Loop 2019–2021
NSF-IIS Eager grant. $300k. PI

Discovering Common Sense from Video, Images, Text and Knowledge Bases 2018–2019
DARPA MCS grant. $640k (my portion). PI of a sub-award for UCLA. PI: Ralph Weischedel (USC)

Learning to Screen: Accelerating Training and Inference for Large NLP Models 2019–2020
Facebook Gift Grant. CO-PI with Cho-Jui Hsieh.

CICI: RDP: Security and Privacy Policy Enforcement for Research Data Protection 2019–2022
NSF-OAC grant. $210k (my portion). PI of a sub-award for UCLA. PI: Yuan Tian (UVirginia)

Online News Trend-Watching via Linguistic Analysis 2018–2023
Taboola Gift fund. Solo-PI

Google GCP Credit Award 2019
$50,000 Google Cloud credits.

Discerning Group Biases in Online Communities via Linguistic Analysis 2018–2019
DARPA UGB grant. $300k (my portion). PI of a sub-award for UCLA. PI: Aram Galstyan (USC)

Reducing Implicit Societal Bias in Artificial Intelligence Systems 2018
Research gift grant, The Okawa Foundataion. Solo-PI

CRII: RI: Learning Structured Prediction Models with Auxiliary Supervision 2016–2019
NSF-IIS grant. $170k. Solo-PI

NVidia GPU Grant 2016
Support 1xTitan X GPU. Solo-PI

TEACHING EXPERIENCE

Instructor, CSM146: Introduction to Machine Learning, UCLA
  - Winter 2018, Fall 2019-2022

Instructor, CS263: Natural Language Processing, UCLA
    Spring 2020-2022

Instructor, CS269: Special Topic in AI: Fairness, Accountability, and Transparency in Natural Language Processing, UCLA

- Winter 2020-2021

Instructor, CS269: Seminar: Machine Learning in Natural Language Processing, UCLA
- Spring 2019, Fall 2017

Instructor, Advanced Machine Learning, University of Virginia
- Spring 2017

Instructor, Natural Language Processing, University of Virginia
- Spring 2017

ADVISEES
Ph.D. Students
- Amita Kamath
- Joe O'Connor
- Joey Yang
- Christina Chance
- Di Wu
- Fan Yin
- Da Yin
- Liunian Harold Li
- Tao Meng
- Anaelia Ovalle
- Arjun Subramonian (Co-advised w/ Sun)
- Tanmay Parekh (Co-advised w/ Peng)
- Hritik Bansal (Co-advised w/ Grover)
- Kareem Ahmad (Co-advised w/ Broeck)
- Pan Lu (Co-advised w/ Zhu) [1]

Alumni
- Kuan-Hao Huang (2017-2023) Current: Postdoc at UIUC
- Md. Rizwan Parvez (2016-2022) Current: Bosch
- Jieyu Zhao (2016-2021) Assistant Professor at USC
- Wasi Uddin Ahmad (2016-2021) Current: Amazon AWS
- Sunipa Dev (2021-2022) Current: Google.
- Dat Duong (Co-advised w/ Eskin, 2017-2020) Current: Scientist at NIH.

TALKS
Tutorials
- Fairness and Biases in Natural Language Processing. **K.-W. Chang**. Machine Learning Summer School 2021.
- Robustness and Adversarial Examples in Natural Language Processing. **K.-W. Chang**, S. Singh, R Jia, H. He. EMNLP 2021.
- Recent Advances in Transferable Representation Learning, M. Chen, **K.-W. Chang**, D. Roth. AAAI 2020.

- Bias and Fairness in Natural Language Processing, **K.-W. Chang**, M. Mitchell and V. Ordonez. EMNLP 2019.
- Quantifying and Reducing Gender Stereotypes in Word Embeddings, **K.-W. Chang**, T. Bolukbasi, and V. Saligrama. FAccT 2018
- Structured Predictions: Practical Advancements and Applications in Natural Language Processing. TAAI 2017.
- Learning and Inference in Structured Prediction Models, **K.-W. Chang**, G. Kundu, D. Roth, and V. Srikumar. AAAI 2016.
- Hands-on Learning to Search for Structured Prediction, H. Daumé III, J. Langford, **K.-W. Chang**, H. He, and S. Rao. NAACL 2015.

Keynote Talks

- NAACL Workshop on Gender Bias in NLP, 2022
- Open Data Science Conference (ODSC) West, 2021.
- West Coast NLP Summit, 2018.
- Southern California Natural Language Processing Symposium, 2017.
- NeurIPS workshop on learning high dimensions with structure, Dec 2016.
- Mid-Atlantic Student Colloquium on Speech, Language and Learning, Breakout Session, Johns Hopkins University, January 2015.

Invited Talks

- Multimodal Representation Learning for Vision and Language
  Baidu 22, Two-Sigma 22
- Unified Pre-Training for Program Understanding and Generation
  Amazon 21, NAACL workshop
- Bias in Language Generation
  Amazon 21.
- Robustness in NLP
  Amazon 20.
- What It Takes to Control Societal Bias in Natural Language Processing
  USC-ISI 20, UCSD 20, UIUC 21, Stanford 18, CMU 18, MSR-NE 19.
- Inject Expert Knowledge and Corpus-Level Constraints in Natural Language Processing Models
  UCLA Stat 19
- Structured Predictions: Practical Advancements and Applications in Natural Language Processing
  UCSD 17, USC 17, Utah 17, Appier Inc. 17, NTU 17, UCLA 17, UMass 17, UCDavis 17
- Multi-Relational Latent Semantic Analysis by Tensor Decomposition
  UMass Med School 16
- Practical Learning Algorithms for Structured Prediction Models
  UMass 15, TTIC 15, UVirginia 15, OSU 15, CMU 15, WSU 15, MSR 15, UArizona 15, MSR-NE 15, Purdue 14, UMD 14, Columbia 14, UIUC 14.

PROFESSIONAL ACTIVITIES

Organizer:

- TrustNLP: Third Workshop on Trustworthy Natural Language Processing. ACL 2023.
- Multilingual Multimodal Learning (MML). ACL 2022.

- TrustNLP: Second Workshop on Trustworthy Natural Language Processing. NAACL 2022.
- Robust and Reliable Machine Learning in the Real World. ICLR 2021.
- TrustNLP: First Workshop on Trustworthy Natural Language Processing. NAACL 2021.
- Workshop on Deep Structured Prediction. ICML 2017.
- 2nd Workshop on Structured Prediction for Natural Language Processing. EMNLP 17.
- Structured Prediction for Natural Language Processing. EMNLP 16.

Associate Program Chair, AAAI 23

Ethics Committee Chair, NAACL 22

Handbook Chair, EMNLP 18

Senior Area Chair: NAACL 21

Area Chair: NLPCC 17, NAACL 18, 19, ACL 19, 20, 21, 22 EMNLP 19, 20, 21, 22 AACL 20,22 ICML 20, ICLR 21, 22, AAAI 20, 21, 22 NeurIPS 22, CoNLL 22, ARR

Senior Program Committee (Area Chair): AAAI 18, 19, IJCAI 20, 21

Journal Editorial Board: JAIR, TACL

Program Committee/Reviewer:

- Machine Learning: ICML (13, 14, 16, 17, 18, 19), ICLR (17), NeurIPS (14, 15, 16, 17, 18, 19), JMLR, TNN, TNNLS, Neurocomputing, Neural Computation.
- Artificial Intelligence: AAAI (14, 15, 16, 17), AISTATS (15, 16), IJPRAI.
- Natural Language Processing: ACL (13, 14, 16, 17, 18), EMNLP (13,15,16,17, 18), NAACL (16, 17), IJCNLP (15,17) CoNLL (14,15,16,17), EACL (17), Coling (14), NLPCC (15), SEM (17).
- Data Mining: KDD (15, 16).

HONORS

Research:

- AAAI Senior Member	2023
- Sloan Research Fellowship	2021
- Best Paper Finalist, CVPR 2022	2022
- Best Long Paper Award, EMNLP 2017	2017
- Best Research Paper Award, SIGKDD 2010	2010
- Okawa Research Grant Award	2018
- Google Research Scholar	2021
- Amazon Research Award	2020
- Facebook Research Award	2020
- NSF CISE Research Initiation Initiative (CRII) Award	2016
- Best Paper Award, ACL Workshop on Representation Learning for NLP	2017
- C. L. and Jane W. S. Liu Award, University of Illinois	2013

  Given in support of a student showing exceptional research promise relatively early in their graduate studies.

- Yahoo! Key Science Challenge Award, Yahoo!	2011
- Master Thesis Award, Taiwanese Association for Artificial Intelligence	2009

- Scholarship for Graduate Student, GARMIN   2008
- Honorary Member of the Phi Tau Phi Scholastic Honor Society, National Taiwan University   2009

Shared Tasks:

- First Place in CoNLL-shared Task 2014   2014
- First Place in CoNLL-shared Task 2013   2013
- Fourth Place in CoNLL-shared Task 2012 English track   2012
- Third Place in CoNLL-shared Task 2011   2011
  CoNLL shared task is the most prestigious competition in NLP.
- Third Prize in the Slow Track of KDDCUP 2009   2009
  Out of 400+ submissions. KDDCup is the most prestigious competition in data mining.