

**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.*
- *FATE: Fairness, Accountability, Transparency, and Ethics in AI.*

**EDUCATION AND EXPERIENCE****University of California Los Angeles, CA**

Assistant Professor, Computer Science

Aug. 2017 –

**University of Virginia, VA**

Assistant Professor, Computer Science

Aug. 2016 – Jun. 2017

**Microsoft, MA**

Postdoctoral researcher, Microsoft Research New England Lab

Jul. 2015 – Aug. 2016

**University of Illinois at Urbana-Champaign, IL**

Ph.D. in Computer Science

Aug. 2010 – May. 2015

*Cognitive Computation Group; w/ Dan Roth*

**National Taiwan University, Taipei, Taiwan**

M.S. in Computer Science and Information Engineering

Jun. 2009

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

B.S. in Computer Science and Information Engineering

Jun. 2007

B.S. in Electrical Engineering (Dual Degree)

Jun. 2007

**Microsoft**

Intern, Microsoft Research, New York (Mentor: John Langford)

2014

Intern, Microsoft Research, Redmond (Mentor: Scott Wen-tau Yih)

2013

Intern, Microsoft Silicon Valley (Mentor: Sathya Keerthi)

2012

**Google**

Intern, Google Research Group, Beijing (Advisor: Chih-Jen Lin)

2008

**SELECTED AWARDS**

Best Long Paper Award, EMNLP 2017

2017

Best Research Paper Award, SIGKDD 2010

2010

Best Paper Award, ACL Workshop on Representation Learning for NLP

2017

Okawa Research Grant Award

2018

NSF CISE Research Initiation Initiative (CRII) Award

2016

Yahoo! Key Science Challenge Award, Yahoo!

2011

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.

## PUBLICATIONS

13,000+ Google Scholar citations in total. H-index: 32, i10-index: 49.

Published at ICML, NeurIPS, EMNLP, NAACL, ACL, AAAI, JMLR, KDD, etc.

### Refereed Conference Publications

- [1] J. Dhamala, T. Sun, V. Kumar, S. Krishna, Y. Pruksachatkun, **K.-W. Chang**, and R. Gupta. BOLD: Dataset and metrics for measuring biases in open-ended language generation. *in FAccT*, 2021.
- [2] 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)*
- [3] Y. Zhou, Y. Yan, R. Han, H. Caufield, **K.-W. Chang**, Y. Sun, P. Ping, and W. Wang. Clinical Temporal Relation Extraction with Probabilistic Soft Logic Regularization and Global Inference. *AAAI Conference on Artificial Intelligence (AAAI 2021)*.
- [4] 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)*.
- [5] 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)*.
- [6] K.-H. Huang, C. Li, and **K.-W. Chang**. Generating Sports News from Live Commentary: A Chinese Dataset for Sports Game Summarization. *Asia Chapter of the Association for Computational Linguistics (ACL 2020)*
- [7] 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)*.
- [8] 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)*.
- [9] 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)*.
- [10] J. Zhao and **K.-W. Chang**. LOGAN: Local Group Bias Detection by Clustering. *Conference on Empirical Methods in Natural Language Processing. (EMNLP 2020, short)*.
- [11] 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)*.
- [12] 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)*
- [13] 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)*
- [14] 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)*
- [15] 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)*
- [16] 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)*
- [17] 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)*
- [18] 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)*
- [19] 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)*
- [20] Z. Shi, H. Zhang, **K.-W. Chang**, M. Huang, and C.-J. Hsieh, Robustness Verification for Transformers. *International Conference on Learning Representation (ICLR 2020)*.
- [21] 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)*.
- [22] 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)*.
- [23] 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)*.
- [24] 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)*.
- [25] 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)*.
- [26] 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)*.
- [27] 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)*.
- [28] 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)*.
- [29] 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)*.
- [30] 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)*.

- [31] 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)*.
- [32] 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)*.
- [33] 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)*.
- [34] 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)*.
- [35] 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)*.
- [36] 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*.
- [37] 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)*.
- [38] 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)*
- [39] 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)*.
- [40] 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)*.
- [41] 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)*.
- [42] 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)*.
- [43] W. Ahmad, **K.-W. Chang**, A Corpus to Learn Refer-to-as Relations for Nominals. *Language Resources and Evaluation Conference (LREC 2018)*.
- [44] 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)*
- [45] 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)*.
- [46] 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)*.
- [47] W. Ahmad, **K.-W. Chang**, H. Wang. Multi-Task Learning for Document Ranking and Query Suggestion. *International Conference on Learning Representation (ICLR 2018)*.

- [48] K. Arnold, **K.-W. Chang**, A. Kalai, Counterfactual Language Model Adaptation for Suggesting Phrases. *International Joint Conference on Natural Language Processing (IJCNLP 2017)*.
- [49] 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**.
- [50] 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)*.
- [51] 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)*.
- [52] 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)*.
- [53] 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)*.
- [54] **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)*.
- [55] 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)*.
- [56] **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)*.
- [57] **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)*.
- [58] R. Samdani, **K.-W. Chang**, D. Roth. A discriminative latent variable model for online clustering. *the 31st International Conference on Machine Learning (ICML 2014)*.
- [59] **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)*.
- [60] **K.-W. Chang**, W.-t. Yih, C. Meek. Multi-relational latent semantic analysis. *Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*.
- [61] **K.-W. Chang**, V. Srikumar, D. Roth. Multi-core structural SVM training. *European Conference on Machine Learning (ECML 2013)*.
- [62] **K.-W. Chang**, S. Sundararajan, S. S. Keerthi. Tractable semi-supervised learning of complex structured prediction models. *European Conference on Machine Learning (ECML 2013)*.
- [63] **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)*.
- [64] **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)*.
- [65] 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)*.
- [66] 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**.

- [67] 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)*.
- [68] 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)*.
- [69] C.-J. Hsieh, **K.-W. Chang**, C.-J. Lin, S. Sathiy 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

- [70] C.-P. Li, **K.-W. Chang**, Distributed Block-diagonal Approximation Methods for Regularized Empirical Risk Minimization, *Machine Learning Journal*, 2019.
- [71] 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
- [72] 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
- [73] 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
- [74] 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.
- [75] 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.
- [76] 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.
- [77] 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.
- [78] **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

- [79] 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)*.
- [80] 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)*.
- [81] 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)*.
- [82] **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).*

- [83] **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).*
- [84] 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

- [85] K.-W. Chang, Selective algorithms for large-scale classification and structured learning (Ph.D.)
- [86] K.-W. Chang, A dual coordinate descent method for large-scale linear SVM (Master)

### Workshop Publications and Preprints

- [87] 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.
- [88] 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.
- [89] W. Shi, M. Chen, Y. Tian, **K.-W. Chang**. Learning Bilingual Word Embeddings Using Lexical Definitions *ACL Representation learning for NLP Workshop 2019*.
- [90] S. Upadhyay, **K.-W. Chang**, M. Taddy, A. Kalai, J. Zou. Beyond Bilingual: Multi-sense Word Embeddings using Multilingual Context. *ACL Representation learning for NLP Workshop 2017, Best Paper Award*.
- [91] K. Arnold, **K.-W. Chang**, A. Kalai. Learning to Suggest Phrases. *AAAI Workshop on Human-Aware AI Workshop*, 2017.
- [92] C.-p. Lee, K.-W. Chang, S. Upadhyay, D. Roth. Distributed Training of Structured SVM. *NeurIPS Workshop on Optimization for Machine Learning*, 2015.
- [93] K. -W. Chang, H. Daumé III, J. Langford, S. Ross. Efficient Programmable Learning to Search. *ICML Workshop on Machine Learning System*, 2015.
- [94] 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 Inferring: Interactions between Inference and Learning*, 2013.
- [95] 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

- [96] 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.
- [97] Interactive Context-Based Text Completions, K. Arnold, **K.-W. Chang**, A. Kalai, 2016 (under review).

### FUNDING

I've secured about \$2,000,000 research fund solely the portion for my group.

<b>Amazon Research Award</b>	2020–2021
<b>AI-DCL: Governing Bias in AI System with Humans in the Decision Loop</b> NSF-IIS Eager grant. \$300k. PI	2019–2021
<b>Discovering Common Sense from Video, Images, Text and Knowledge Bases</b> DARPA MCS grant. \$640k (my portion). PI of a sub-award for UCLA. PI: Ralph Weischedel (USC)	2018–2019
<b>Learning to Screen: Accelerating Training and Inference for Large NLP Models</b> Facebook Gift Grant. CO-PI with Cho-Jui Hsieh.	2019–2020
<b>CICI: RDP: Security and Privacy Policy Enforcement for Research Data Protection</b> NSF-OAC grant. \$210k (my portion). PI of a sub-award for UCLA. PI: Yuan Tian (UVirginia)	2019–2022
<b>Online News Trend-Watching via Linguistic Analysis</b> Taboola Gift fund. Solo-PI	2018–2021
<b>Google GCP Credit Award</b> \$50,000 Google Cloud credits.	2019
<b>Discerning Group Biases in Online Communities via Linguistic Analysis</b> DARPA UGB grant. \$300k (my portion). PI of a sub-award for UCLA. PI: Aram Galstyan (USC)	2018–2019
<b>Reducing Implicit Societal Bias in Artificial Intelligence Systems</b> Research gift grant, The Okawa Foundataion. Solo-PI	2018
<b>CRII: RI: Learning Structured Prediction Models with Auxiliary Supervision</b> NSF-IIS grant. \$170k. Solo-PI	2016–2019
<b>NVidia GPU Grant</b> Support 1xTitan X GPU. Solo-PI	2016

## TEACHING EXPERIENCE

Instructor, CSM146: Introduction to Machine Learning, UCLA

- Fall 2020, enrollment: 220
- Fall 2019, enrollment: 191
- Fall 2018, enrollment: 155
- Winter 2018, enrollment: 193

Instructor, CS263: Natural Language Processing, UCLA

- Spring 2020, enrollment: 76

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

- Winter 2020, enrollment: 39

Instructor, CS269: Seminar: Machine Learning in Natural Language Processing, UCLA

- Spring 2019, enrollment: 38
- Fall 2017, enrollment: 76

Instructor, Advanced Machine Learning, University of Virginia

- Spring 2017, enrollment:  $\approx 50$

Instructor, Natural Language Processing, University of Virginia

- Spring 2017, enrollment:  $\approx 50$



## ADVISEES

### Ph.D. Students

- Fan Yin (expected graduation ~2025)
- Da Yin (expected graduation ~2025)
- Liunian Harold Li (expected graduation ~2024)
- Kareem Ahmad (Co-advised w/ Broeck, expected graduation ~2024)
- Kuan-Hao Huang (expected graduation ~2023)
- Tao Meng (expected graduation ~2023)
- Wasi Uddin Ahmad (expected graduation ~2022)
- Jieyu Zhao (expected graduation ~2022)
- Md. Rizwan Parvez (expected graduation ~2022)
- Dat Duong (Co-advised w/ Eskin, expected graduation ~2020)

### Undergraduate and M.S. Students

- Shanxiu He (expected graduation ~2021)
- Yihe Deng (expected graduation ~2021)
- Di Qu (expected graduation ~2021)
- Jesse Cai (expected graduation ~2020)
- Zhubo Deng (expected graduation ~2020)
- Weijia Shi (expected graduation ~2020)

### Past Undergraduate and M.S. Students

- **Summer Students:** Da Yin, Eleanor Jiang, Fan Yin, Quanyu Long, Shengyu Jia, Zhenxin Xiao, Siyu Wang, Junyu Guo, Hanwen Jiang
- **M.S. students:** Haoxiang Zhang, Xia Chen, Sudharsan Krishnaswamy, Puchin Chen (Google), Peter Kim, Evelyn Chen (Amazon), Cindy Chiang, Daphne Chiou (Google)
- **Undergraduate students:** Chao Jiang (PhD at OSU), Zhechao Huang (MS at MIT)

## TALKS

### Tutorials

- 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. FAT 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

- 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

- What It Takes to Control Societal Bias in Natural Language Processing  
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:

- Handbook Chair, 2018
- Workshop on Deep Structured Prediction. Co-organizers: Isabelle Augenstein, Kai-Wei Chang, Gal Chechik, Bert Huang, Andre Martins, Ofer Meshi, Yishu Miao, Alexander Schwing. ICML 2017.
- 2nd Workshop on Structured Prediction for Natural Language Processing. Co-organizers: K.-W. Chang, M. Chang, A. Rush, and V. Srikumar. EMNLP 2017.
- Structured Prediction for Natural Language Processing. Co-organizers: K.-W. Chang, M. Chang, A. Rush, and V. Srikumar. EMNLP 2016.

Senior Program Committee / Area Chair: NLPCC 17, AAI 18, 19, NAACL 18, 19, ACL 19, 20 EMNLP 19, 20, ACL 20, ICML 20, AAI 20, 21 (Area Chair).

### Program Committee/Reviewer/Editorial Board:

- 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: AAI (14, 15, 16, 17), AISTATS (15, 16), IJPRAI, JAIR.
- 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), EAACL (17), Coling (14), NLPCC (15), SEM (17), TACL.
- Data Mining: KDD (15, 16).

## OTHER HONORS

### Research:

- 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
  - Facebook PhD Fellowship Finalist 2014
  - Microsoft Research PhD Fellowship Finalist 2013
  - Studying Abroad Scholarship, Ministry of Education, Taiwan 2009  
Given to top 3% master students.
  - Undergraduate Research Grant, National Science Council, Taiwan 2006

**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.**