

TONGFEI CHEN

✉ tongfeichen@microsoft.com ✉ tongfei@pm.me 🌐 tongfei.me

EDUCATION

- 2020 | **Ph.D. in Computer Science**
Johns Hopkins University, Baltimore, MD, USA
Center for Language and Speech Processing
Department of Computer Science
Advisor: Prof. Benjamin Van Durme
Thesis: *Ranking and retrieval under semantic relevance*
- 2014 | **B.Sc. in Computer Science**
Peking University, Beijing, China
Department of Computer Science and Technology
Advisor: Prof. Junfeng Hu
Thesis: *Large-scale unsupervised word segmentation for Classical Chinese: Research & system*

WORK EXPERIENCE

- Jan. 2022 – present | **Senior Researcher**
Semantic Machines, Microsoft, Bellevue / Redmond, WA, USA
Topics: Conversational AI; dialogue systems; few-shot semantic parsing; dynamic prompt creation
- Oct. 2020 – Dec. 2021 | **Senior Applied Scientist**
Project Turing, Microsoft, Bellevue, WA, USA
Topics: Microsoft Bing; Transformer model pre-training; question answering over tables; relevance ranking for tables
- Aug. 2014 – Sept. 2020 | **Graduate Research Assistant**
Center for Language and Speech Processing, Johns Hopkins University, Baltimore, MD, USA
Advisor: Prof. Benjamin Van Durme
Topics: Information extraction; question answering; information retrieval; natural language inference; knowledge acquisition from text; scalable systems; approximate algorithms
- May 2018 – Aug. 2018 | **Applied Scientist Intern**
Alexa AI, Amazon.com, Seattle, WA, USA
Host: Dr. Lambert Mathias
Topics: Amazon Alexa; dialogue context modeling; sequence transduction models
- May 2017 – Aug. 2017 | **Research Intern**
IBM Thomas J. Watson Research Center, IBM, Yorktown Heights, NY, USA
Hosts: Dr. Jiří Navrátil
Topics: Confidence scoring; model calibration; meta-models
- June 2012 – June 2014 | **Undergraduate Research Assistant**
Institute of Computational Linguistics, Peking University, Beijing, China
Advisor: Prof. Junfeng Hu
Topics: Word segmentation; graph & network analysis; ontology construction

PUBLICATIONS

Preprints and working papers

- [P1] Subhro Roy, Sam Thomson, [Tongfei Chen](#), Richard Shin, Adam Pauls, Jason Eisner, Benjamin Van Durme (2022). [BenchCLAMP: A benchmark for evaluating language models on semantic parsing](#). [NLP]

Peer-reviewed papers

* Equal contributions.

- [1] Yunmo Chen, William Gantt, Weiwei Gu, [Tongfei Chen](#), Aaron Steven White, Benjamin Van Durme (2023). [Iterative document-level information extraction via imitation learning](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. pp. 1858–1874. **[outstanding paper award]** [NLP]
- [2] Weiwei Gu*, Boyuan Zheng*, Yunmo Chen, [Tongfei Chen](#), Benjamin Van Durme (2022). [An empirical study on finding spans](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pp. 3976–3983. [NLP]
- [3] Luyu Gao, Zhuyun Dai, [Tongfei Chen](#), Zhen Fan, Benjamin Van Durme, Jamie Callan (2021). [Complement lexical retrieval model with residual semantic embeddings](#). In *Proceedings of the 43rd European Conference on Information Retrieval (ECIR; LNCS 12656)*. pp. 146–160. [IR] [NLP]
- [4] Patrick Xia*, Guanghui Qin*, Siddharth Vashishtha, Yunmo Chen, [Tongfei Chen](#), Chandler May, Craig Harman, Kyle Rawlins, Aaron Steven White, Benjamin Van Durme (2021). [LOME: Large ontology multilingual extraction](#). In *Proceedings of the Software Demonstrations of the 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. pp. 149–159. [NLP]
- [5] Yunmo Chen, [Tongfei Chen](#), Benjamin Van Durme (2020). [Joint modeling of arguments for event understanding](#). In *Proceedings of the First Workshop on Computational Approaches to Discourse (CODI@EMNLP)*. pp. 96–101. [NLP]
- [6] Yunmo Chen, [Tongfei Chen](#), Seth Ebner, Aaron Steven White, Benjamin Van Durme (2020). [Reading the Manual: Event extraction as definition comprehension](#). In *Proceedings of the Fourth Workshop on Structured Prediction for NLP (SPNLP@EMNLP)*. pp. 74–83. [NLP]
- [7] [Tongfei Chen](#)*, Zhengping Jiang*, Adam Poliak, Keisuke Sakaguchi, Benjamin Van Durme (2020). [Uncertain natural language inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. pp. 8772–8779. [NLP]
- [8] [Tongfei Chen](#), Yunmo Chen, Benjamin Van Durme (2020). [Hierarchical entity typing via multi-level learning to rank](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. pp. 8465–8475. [NLP]
- [9] Yiming Wang, [Tongfei Chen](#), Hainan Xu, Shuoyang Ding, Hang Lv, Yiwen Shao, Nanyun Peng, Lei Xie, Shinji Watanabe, Sanjeev Khudanpur (2019). [ESPRESSO: A fast end-to-end neural speech recognition toolkit](#). In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. pp. 136–143. [SPEECH]
- [10] Arya D. McCarthy, [Tongfei Chen](#), Seth Ebner (2019). [An exact No Free Lunch theorem for community detection](#). In *Proceedings of the 8th International Conference on Complex Networks and their Applications (COMPLEX NETWORKS)*. pp. 176–187. [GRAPH]
- [11] Arya D. McCarthy, [Tongfei Chen](#), Rachel Rudinger, David W. Matula (2019). [Metrics matter in community detection](#). In *Proceedings of the 8th International Conference on Complex Networks and their Applications (COMPLEX NETWORKS)*. pp. 164–175. [GRAPH]
- [12] [Tongfei Chen](#), Chetan Naik, Hua He, Pushpendre Rastogi, Lambert Mathias (2019). [Improving long distance slot carryover in spoken dialogue systems](#). In *Proceedings of the First Workshop of NLP for Conversational AI (NLP4ConvAI@ACL)*. pp. 96–105. **[best paper award]** [NLP]

- [13] Zhongyang Li, [Tongfei Chen](#), Benjamin Van Durme (2019). [Learning to rank for plausible plausibility](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*. pp. 4818–4823. [NLP]
- [14] Pushpendre Rastogi, Arpit Gupta, [Tongfei Chen](#), Lambert Mathias (2019). [Scaling multi-domain dialogue state tracking via query reformulation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (NAACL)*. pp. 97–105. [NLP]
- [15] J. Edward Hu, Huda Khayrallah, Ryan Culkin, Patrick Xia, [Tongfei Chen](#), Matt Post, Benjamin Van Durme (2019). [Improved lexically constrained decoding for translation and monolingual rewriting](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (NAACL)*. pp. 839–850. [NLP]
- [16] Yiming Wang, Xing Fan, I-Fan Chen, Yuzong Liu, [Tongfei Chen](#), Björn Hoffmeister (2019). [End-to-end anchored speech recognition](#). In *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 7090–7094. [SPEECH]
- [17] [Tongfei Chen](#), Jiří Navrátil, Vijay Iyengar, Karthikeyan Shanmugam (2019). [Confidence scoring using whitebox meta-models with linear classifier probes](#). In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS); Proceedings of Machine Learning Research 89 (PMLR)*, pp. 1467–1475. [ML]
- [18] Rashmi Sankepally, [Tongfei Chen](#), Benjamin Van Durme, Douglas W. Oard (2018). [A test collection for coreferent mention retrieval](#). In *Proceedings of the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pp. 1209–1212. [NLP]
- [19] Hainan Xu, [Tongfei Chen](#), Dongji Gao, Yiming Wang, Ke Li, Nagendra Goel, Yishay Carmiel, Daniel Povey, Sanjeev Khudanpur (2018). [A pruned RNNLM lattice-rescoring algorithm for automatic speech recognition](#). In *Proceedings of the 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5929–5933. [SPEECH] [NLP]
- [20] [Tongfei Chen](#) (2017). [Typesafe abstractions for tensor operations](#). In *Proceedings of the 8th ACM SIGPLAN International Symposium on Scala (SCALA@SPLASH)*. pp. 45–50. [PL]
- [21] Benjamin Van Durme, Tom Lippincott, Kevin Duh, Deana Burchfield, Adam Poliak, Cash Costello, Tim Finin, Scott Miller, James Mayfield, Philipp Koehn, Craig Harman, Dawn Lawrie, Chandler May, Max Thomas, Annabelle Carrell, Julianne Chaloux, [Tongfei Chen](#), Alex Comerford, Mark Dredze, Benjamin Glass, Shudong Hao, Patrick Martin, Pushpendre Rastogi, Rashmi Sankepally, Travis Wolfe, Ying-Ying Tran, Ted Zhang (2017). [CADET: Computer Assisted Discovery Extraction and Translation](#). In *Proceedings of the 8th International Joint Conference on Natural Language Processing, System Demonstrations (IJCNLP)*, pp. 5–8. [NLP]
- [22] [Tongfei Chen](#), Benjamin Van Durme (2017). [Discriminative information retrieval for question answering sentence selection](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2 (EACL)*, pp. 719–725. [NLP] [IR]
- [23] Junhao Zhang, [Tongfei Chen](#), Junfeng Hu (2015). [On the relationship between Gaussian stochastic blockmodels and label propagation algorithms](#). *Journal of Statistical Mechanics: Theory and Experiment (J. Stat. Mech)*. 2015(3), P03009. [GRAPH]
- [24] Ni Sun, [Tongfei Chen](#), Liumingjing Xiao, Junfeng Hu (2014). [Diachronic deviation features in continuous space word representations](#). In *Proceedings of the 13th China National Conference on Computational Linguistics (CCL; LNCS 8801)*, pp. 23–33. [NLP]
- [25] [Tongfei Chen](#), Xiaojun Zou, Weimeng Zhu, Junfeng Hu (2013). [Human-computer interactive Chinese word segmentation: An adaptive Dirichlet process mixture model approach](#). In *Proceedings of the 6th International Joint Conference on Natural Language Processing (IJCNLP)*, pp. 1278–1284. [NLP]

- [26] **Tongfei Chen**, Weimeng Zhu, Xueqiang Lv, Junfeng Hu (2013). *A Kalman filter based human-computer interactive segmentation system for ancient Chinese texts*. In *Proceedings of the 12th China National Conference on Computational Linguistics (CCL; LNCS 8202)*, pp. 25–35. [NLP]

Non-refereed system descriptions

- [S1] Yunmo Chen, Seth Ebner, **Tongfei Chen**, Patrick Xia, Elias Stengel-Eskin, Tzu-Ray Su, J. Edward Hu, Nils Holzenberger, Ryan Culkin, Craig Harman, Max Thomas, Thomas Lippincott, Aaron Steven White, Kyle Rawlins, Benjamin Van Durme (2019). *NIST TAC SM-KBP 2019 system description: JHU/UR framework*. In *Proceedings of the Text Analysis Conference (TAC)*. [NLP]
- [S2] Mozhi Zhang, Jordan Boyd-Graber, Michelle Yuan, C. Anton Rytting, Weiwei Yang, Philip Resnik, Ting Hua, Adam Poliak, Adam Teichert, **Tongfei Chen**, Xu Han, Linghao Jin, João Sedoc, Benjamin Van Durme (2019). *LoReHLT19 System Description UMD-JHU*. [NLP]
- [S3] Patrick Xia, Elias Stengel-Eskin, **Tongfei Chen**, Seth Ebner, Nils Holzenberger, Ryan Culkin, Pushpendre Rastogi, Xutai Ma, Benjamin Van Durme (2018). *NIST TAC SM-KBP 2018 system description: JHU/UR pipeline*. In *Proceedings of the Text Analysis Conference (TAC)*. [NLP]

HONORS AND AWARDS

- | | |
|-----------|--|
| May 2023 | Outstanding paper award (10/281)
<i>The 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL)</i> |
| July 2019 | Best paper award (2/25)
<i>The 1st Workshop on NLP for Conversational AI @ ACL 2019</i> |
| June 2014 | Outstanding undergraduate thesis
<i>Peking University</i> |

PRESENTATION AND TALKS

- | | |
|---------------|---|
| Jul. 8, 2020 | Hierarchical entity typing via multi-level learning to rank
<i>Annual Meeting of the Association for Computational Linguistics (ACL) 2020, Virtual online</i> |
| Jul. 8, 2020 | Uncertain natural language inference
<i>Annual Meeting of the Association for Computational Linguistics (ACL) 2020, Virtual online</i> |
| Aug. 1, 2019 | Improving long distance slot carryover in spoken dialogue systems
<i>Workshop on NLP for Conversational AI @ ACL 2019: Best paper talk, Florence, Tuscany, Italy</i> |
| May 29, 2019 | Uncertain natural language inference
<i>DARPA LORELEI/AIDA Site Visit, JHU</i> |
| Nov. 16, 2018 | Towards typesafe deep learning in Scala
<i>Scale by the Bay 2018, San Francisco, CA, USA</i> |
| Mar. 18, 2018 | Towards typesafe deep learning in Scala
<i>Northeast Scala Symposium 2018, Cambridge, MA, USA</i> |
| Oct. 23, 2017 | Typesafe abstractions for tensor operations
<i>SCALA@SPLASH 2017, Vancouver, BC, Canada</i> |

- Feb. 15, 2017 | Discriminative information retrieval for knowledge discovery
DARPA DEFT/LORELEI Site Visit, JHU
- Oct. 17, 2013 | HCI Chinese word segmentation: An adaptive Dirichlet process mixture model approach
International Joint Conference on Natural Language Processing (IJCNLP) 2013, Nagoya, Aichi, Japan
- Oct. 11, 2013 | A Kalman filter based HCI segmentation system for ancient Chinese texts
Chinese Computational Linguistics (CCL) 2013, Suzhou, Jiangsu, China

OPEN SOURCE

- Espresso: Fast end-to-end automatic speech recognition based on fairseq. (800+ stars on GitHub)
<https://github.com/freewym/espresso> (contributor)
- Nexus: Experimental typesafe tensors and deep learning in Scala. (250+ stars on GitHub)
<https://github.com/ctongfei/nexus>
- Progressbar: A terminal-based progress bar for JVM. (900+ stars on GitHub)
<https://github.com/ctongfei/progressbar>

SERVICE

- Meta-reviewer / Senior program committee member:
 - AAAI 2022
- Reviewer / Program committee member:

<ul style="list-style-type: none"> – ACL Rolling Review 2022 – AAAI 2021, 2020 – AACL 2020 – ACL 2023, 2021, 2020 – CCL 2017 – <i>Computational Linguistics</i> 2023, 2022, 2021 – DI@KDD 2021 	<ul style="list-style-type: none"> – EACL 2023 – EMNLP 2022, 2021, 2020, 2019, 2018 – KG4IR@SIGIR 2017 – NAACL 2021, 2019 – RCQA@AAAI 2020 – TADGM@ICML 2018
---	--
- Secondary reviewer:

<ul style="list-style-type: none"> – ACL 2019, 2018, 2017, 2015, 2014 – ACL Demo Track 2017 – EACL 2017 – EMNLP 2017, 2014 	<ul style="list-style-type: none"> – IJCNLP 2017 – NAACL 2015 – TACL 2017, 2015 – WWW 2015
--	--
- PhD recruitment committee 2018–2020, Johns Hopkins University
- North American Computational Linguistics Olympiad (NACLO) organizing committee 2016

SKILLS

- Programming languages:

<ul style="list-style-type: none"> – Scala (<i>expert</i>) – Python (<i>proficient</i>) – Java (<i>proficient</i>) 	<ul style="list-style-type: none"> – C/C++ (<i>proficient</i>) – C# (<i>proficient</i>) – Haskell (<i>intermediate</i>)
---	--
- Natural languages:

- Mandarin Chinese (*native*)
 - English (*proficient*)
 - Japanese (*intermediate*)
- Libraries and tools:
 - Deep learning: PyTorch, TensorFlow
 - Information retrieval: Lucene, FAISS
 - Data serialization: Thrift, ProtoBuf
 - Data visualization: Gephi
 - Distributed computing: Spark
 - Workflow orchestration: Ducttape
 - Scala ecosystem: Cats, Shapeless
 - Typesetting: \LaTeX