

Jason Wei

jason.weng.wei@gmail.com
https://jasonwei20.github.io

Employment

2020 - *Google Brain*
Research Scientist, Jun. 2022 - present.
Research Engineer, Dec. 2021 - Jun. 2022.
AI Resident, Oct. 2020 - Dec. 2021.

Education

2016 - 2020 *Dartmouth College*
Bachelor of Arts, Computer Science.
Thesis Advisor: Lorenzo Torresani.

Publications

Emergent abilities of large language models.

J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, and W. Fedus
arXiv preprint.

Least-to-most prompting enables complex reasoning in large language models.

D. Zhou, N. Schärli, L. Hou, J. Wei, N. Scales, X. Wang, D. Schuurmans, O. Bousquet, Q. Le, and E. Chi.
arXiv preprint.

PaLM: Scaling language modeling with Pathways.

{A. Chowdhery, S. Narang, J. Devlin}, et al. (64 additional authors).
arXiv preprint.

Self-consistency improves chain of thought reasoning in language models.

X. Wang, J. Wei, D. Schuurmans, Q. Le, E. Chi, S. Narang, A. Chowdhery, and D. Zhou.
arXiv preprint.

Chain of thought prompting elicits reasoning in large language models.

J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. Chi, Q. Le, and D. Zhou.
arXiv preprint.

ACL '22 *A recipe for arbitrary text style transfer with large language models.*

{Emily Reif, Daphne Ippolito}, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei.
Conference of the Association for Computational Linguistics, 2022.

ICLR '22 *Finetuned language models are zero-shot learners.*

{Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu}, Adams Wei Yu, Brian Lester, Nan Du, Andrew Dai, and Quoc Le.
International Conference on Learning Representations, 2022 (oral).

ICLR '22 *The MultiBERTs: BERT reproductions for robustness analysis.*

{Thibault Sellam, Steve Yadlowsky}, Ian Tenney, Jason Wei, Naomi Saphra, Alexander D'Amour, Tal Linzen, Jasmijn Bastings, Iulia Turc, Jacob Eisenstein, Dipanjan Das, and Ellie Pavlick.
International Conference on Learning Representations, 2022.

- EMNLP '21 *Frequency effects on syntactic rule-learning in transformers.*
Jason Wei, Dan Garrette, Tal Linzen, and Ellie Pavlick.
Conference on Empirical Methods in Natural Language Processing, 2021 (oral).
- EMNLP '21 *Good-enough example extrapolation.*
Jason Wei.
Conference on Empirical Methods in Natural Language Processing, 2021.
- ACL '21 *A cognitive regularizer for language modeling.*
Jason Wei, Clara Meister, and Ryan Cotterell.
Conference of the Association for Computational Linguistics, 2021.
- ACL '21 *Language model augmented relevance score.*
Ruibo Liu, Jason Wei, and Soroush Vosoughi.
Conference of the Association for Computational Linguistics, 2021 (oral).
- ACL '21 (Findings) *A survey of data augmentation approaches for NLP.*
{S. Feng, V. Gangal}, J. Wei, S. Chandar, S. Vosoughi, T. Mitamura, and E. Hovy.
Findings of the Association for Computational Linguistics: ACL 2021.
- ACL '21 (Findings) *Modulating language models with emotions.*
Ruibo Liu, Jason Wei, Chenyan Jia, and Soroush Vosoughi.
Findings of the Association for Computational Linguistics: ACL 2021.
- NAACL '21 *Linguistic complexity loss in text-based therapy.*
Jason Wei, Kelly Finn, Emma Templeton, Thalia Wheatley, and Soroush Vosoughi.
Conference of the North American Chapter of the Association for Computational Linguistics, 2021.
- NAACL '21 *Few-shot text classification with triplet networks, data augmentation, and curriculum learning.*
Jason Wei, Chengyu Huang, Soroush Vosoughi, Yu Cheng, and Shiqi Xu.
Conference of the North American Chapter of the Association for Computational Linguistics, 2021.
- EACL '21 *Text augmentation in a multi-task view.*
Jason Wei, Chengyu Huang, Shiqi Xu, and Soroush Vosoughi.
Conference of the European Chapter of the Association for Computational Linguistics, 2021.
- AAAI '21 *Mitigating political bias in language models through reinforced calibration.*
Ruibo Liu, Chenyan Jia, Jason Wei, Guangxuan Xu, Lili Wang, and Soroush Vosoughi.
AAAI Conference on Artificial Intelligence, 2021 (Outstanding AI for Social Impact Paper).
- EMNLP '19 *Easy data augmentation techniques for boosting performance on text classification tasks.*
Jason Wei and Kai Zou.
Conference on Empirical Methods in Natural Language Processing, 2019.

Invited Talks

- 2022 "Emergent abilities of big language models."
- Princeton NLP group, Feb. 2022.
- 2022 "What is natural language processing and why care?"
- Dartmouth College, Feb. 2022.
- 2022 "Instruction tuning with FLAN."
- Stanford NLP Seminar, Jan. 2022.
- New York University, ML² group, Dec. 2021.
- Google Research, NLP Discuss Series, Nov. 2021.
- 2021 "When BERT fails at syntax."
- New York University, Computation and Psycholinguistics Lab, Nov. 2021.

Honors

- 2021 AAAI-21 AI for Social Impact Outstanding Paper (Liu et al.).
- 2020 Phi Beta Kappa, Tau Beta Pi, Gamma Sigma Alpha.
- 2020 Neukom Prize for Outstanding Undergraduate Research.
- 2019 Barry Goldwater Scholarship.
- 2016 Intel Science Talent Search Semifinalist.

Teaching

Machine Learning and Statistical Data Analysis (COSC 74/174).
Teaching Assistant, Dartmouth College, Spring 2018.

Service

Reviewing: ICML '22, ICLR '22, ACL '21, EMNLP '21, NAACL '21, EACL '21, ACL Rolling Review '21.