Peng Qian

Address
Department of Psychology
Harvard University
Northwest Building Room 180.03
52 Oxford St., Cambridge, MA 02138 USA

Contact
pqian@fas.harvard.edu
pqian@mit.edu
https://www.mit.edu/~pqian

Appointment

Postdoctoral Fellow
Department of Psychology, Harvard University
Department of Brain and Cognitive Sciences, MIT — 2022 - current

Education

Ph.D. Cognitive Science
Department of Brain and Cognitive Sciences, MIT — 2017 - 2022

M.Eng. Intelligent Media Computing
School of Computer Science, Fudan University — 2015 - 2017

B.A. Chinese Language
Department of Chinese Language and Literature, Fudan University — 2011 - 2015

IASP Exchange Programme
Morningside College, Chinese University of Hong Kong — 2013

Preprints (⋆: equal contribution)

1. **Qian, P.**, & Ullman, T. Shape guides visual pretense. *PsyArXiv*.
2. **Qian, P.**, & Levy, R. P. Comprehenders' error correction mechanisms are finely calibrated to language production statistics. *PsyArXiv*.
3. Bridgers, S. E. C.⋆, **Qian, P.**⋆, Parece, K., Taliaferro, M., Schulz, L., & Ullman, T. Loopholes: A window into value alignment and the communication of meaning. *PsyArXiv*.
4. Parece, K., Bridgers, S. E. C., **Qian, P.**, Schulz, L., & Ullman, T. Skirting the sacred: Moral contexts increase the cost of intentional misunderstandings. *PsyArXiv*.
5. Xu, N., Zhang, Q., Zhang, M., **Qian, P.**⋆, & Huang, X⋆. On the tip of the tongue: Analyzing conceptual representation in large language models with reverse-dictionary probe. *arXiv preprint arXiv:2402.14404*.

Publications

1. **Qian, P.**, Bridgers, S. E. C., Taliaferro, M., Parece, K., & Ullman, T. (*in press*). Ambivalence by design: A computational account of loopholes. *Cognition*.
2. Murthy, S., Parece, K., Bridgers, S., **Qian, P.**, & Ullman, T. (2023). Comparing the evaluation and production of loophole behavior in humans and large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023* (pp. 4010-4025).
3. **Qian, P.**, & Levy, R. (2022). Flexible generation from fragmentary linguistic input. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 8176-8196).
4. Wilcox, E. G., Gauthier, J., Hu, J., **Qian, P.**, & Levy, R. (2022). Learning syntactic structures from string input. In *Algebraic Structures in Natural Language* (pp. 113-138). CRC Press.

5. Tucker, M., Eisape, T., **Qian, P.**, Levy, R., & Shah, J. (2022). When does syntax mediate neural language model performance? Evidence from dropout probes. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 5393-5408).

6. Wang, Y., Hu, J., Levy, R., & **Qian, P.** (2021). Controlled evaluation of grammatical knowledge in Mandarin Chinese language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing* (pp. 5604-5620).

7. **Qian, P.**, Naseem, T., Levy, R., & Astudillo, R. F. (2021). Structural guidance for Transformer language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (pp. 3735-3745).

8. Tucker, M., **Qian, P.**, & Levy, R. (2021). What if this modified that? Syntactic interventions with counterfactual embeddings. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021* (pp. 862-875).

9. Wilcox, E., **Qian, P.**, Futrell, R., Kohita, R., Levy, R., & Ballesteros, M. (2020). Structural supervision improves few-shot learning and syntactic generalization in neural language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 4640-4652).

10. Wilcox, E. G., Gauthier, J., Hu, J., **Qian, P.**, & Levy, R. P. (2020). On the predictive power of neural language models for human real-time comprehension behavior. In *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society (CogSci), 2020.*

11. Gauthier, J., Hu, J., Wilcox, E., **Qian, P.**, & Levy, R. (2020). SyntaxGym: An online platform for targeted evaluation of language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations* (pp. 70-76).

12. Hu, J., Gauthier, J., **Qian, P.**, Wilcox, E., & Levy, R. (2020). A systematic assessment of syntactic generalization in neural language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 1725-1744).

13. Mollica, F., Siegelman, M., Diachek, E., Piantadosi, S. T., Mineroff, Z., Futrell, R., Kean, H., **Qian, P.**, & Fedorenko, E. (2020). Composition is the core driver of the language-selective network. *Neurobiology of Language*, *1*(1), 104-134.

14. An, A., **Qian, P.**, Wilcox, E., & Levy, R. (2019). Representation of constituents in neural language models: Coordination phrase as a case study. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 2888-2899).

15. **Qian, P.**, Hewitt, L., Tenenbaum, J. B., & Levy, R. (2019). Inferring structured visual concepts from minimal data. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 41).

16. Futrell, R., **Qian, P.**, Gibson, E., Fedorenko, E., & Blank, I. (2019). Syntactic dependencies correspond to word pairs with high mutual information. In *Proceedings of the fifth international conference on dependency linguistics (DepLing, syntaxfest 2019)* (pp. 3-13).

17. Wilcox, E., **Qian, P.**, Futrell, R., Ballesteros, M., & Levy, R. (2019). Structural supervision improves learning of non-local grammatical dependencies. In *Proceedings of NAACL-HLT* (pp. 3302-3312).

18. Futrell, R., Wilcox, E., Morita, T., **Qian, P.**, Ballesteros, M., & Levy, R. (2019). Neural language models as psycholinguistic subjects: Representations of syntactic state. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (pp. 32-42).

19. **Qian, P.**, Qiu, X., & Huang, X. (2016). Analyzing linguistic knowledge in sequential model of sentence. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 826-835).

20. **Qian, P.**, Qiu, X., & Huang, X. (2016). Investigating language universal and specific properties in word embeddings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1478-1488).

21. **Qian, P.**, Qiu, X., & Huang, X. (2016). A new psychometric-inspired evaluation metric for Chinese word segmentation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 2185-2194).

22. **Qian, P.**, Qiu, X., & Huang, X. (2016). Bridging LSTM architecture and the neural dynamics during reading. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence* (pp. 1953-1959).

23. Qiu, X., **Qian, P.**, & Shi, Z. (2016). Overview of the NLPCC-ICCPOL 2016 shared task: Chinese word segmentation for micro-blog texts. In *Natural Language Understanding and Intelligent Applications*, (pp. 901-906). Springer International Publishing

24. Qiu, X., **Qian, P.**, Yin, L., Wu, S., & Huang, X. (2015). Overview of the NLPCC 2015 shared task: Chinese word segmentation and POS tagging for micro-blog texts. In *Natural Language Processing and Chinese Computing: 4th CCF Conference, NLPCC 2015, Nanchang, China, October 9-13, 2015, Proceedings 4* (pp. 541-549). Springer International Publishing.

25. **Qian, P.**, Huang, X. Statistical modelling and macro-analysis of Chinese classical poetry. (2015). *Journal of Jiangxi Normal University (Natural Sciences Edition)*, 2015(2).

26. Chen, Z., **Qian, P.**. Entailment, presupposition, and integrity. (2015). *Contemporary Linguistics*, 2015(1):98-109.

27. **Qian, P.** (2013). Gesture might have laid foundations for the emergence of spoken languages. *Communication on Contemporary Anthropology*, 2013(7):114-118.

Extended Abstracts

1. **Qian, P.**, Parece, K., & Ullman, T. (2024). The benefit of the doubt: People expect greater punishment reduction than predicted by naive plausible deniability. *The 50th Annual Meeting of the Society for Philosophy and Psychology.*

2. **Qian, P.**, & Ullman, T. (2023). Human visual pretense is constrained primarily by shape. *The 49th Annual Meeting of the Society for Philosophy and Psychology.*

3. **Qian, P.**, & Levy, R. (2022). Noisy-channel inference in correcting English number agreement mismatch. *The 35th Annual Conference on Human Sentence Processing.*

4. Wang, Y., Hu, J., Levy, R., & **Qian, P**. (2022). Facilitative effect induced by classifier-noun mismatch in Mandarin Chinese. *The 35th Annual Conference on Human Sentence Processing.*
5. **Qian, P.**, & Levy, R. (2021). Flexibility in language production: Insights from completion of fragmentary inputs. *The 34th Annual Conference on Human Sentence Processing.*
6. Wilcox, E., Gauthier, J., Hu, J., **Qian, P.**, & Levy, R. (2020). Benchmarking neural networks as models of human language processing. *Proceedings of the 26th Architectures and Mechanisms for Language Processing Conference.*
7. Wilcox, E., Gauthier, J., **Qian, P.**, Hu, J., & Levy, R. (2020). Evaluating the effect of model inductive bias and training data in predicting human reading times. *Proceedings of the 33rd Annual CUNY Human Sentence Processing Conference.*
8. Wilcox, E., **Qian, P.**, Futrell, R., Ballesteros, M., & Levy, R. (2019). Syntactic structure aids learning of grammatical dependencies in neural networks. *The 32nd Annual CUNY Conference on Sentence Processing.*
9. Futrell, R., Wilcox, E., Morita, T., **Qian, P.**, Ballesteros, M., & Levy, R. (2019). Large neural network language models learn representations of incremental parse states. *The 32nd Annual CUNY Conference on Sentence Processing.*

Teaching

Teaching fellow at the Workshop for Computational Cognitive Models of Learning and Development	2023
Certificate of Graduate Teaching Development Track	Sept 2021
Completion of Inclusive Teaching Track	June 2021
Completion of Microteaching	April 2021
Completion of Lesson Planning Track	Jan 2021
Completion of Subject Design Track	Dec 2020
Teaching assistant for 9.66J/9.660/6.804J Computational Cognitive Science	2019
Teaching assistant for 9.19/9.190 Computational Psycholinguistics	2019

Service

Reviewer for *Open Mind* (2023, 2024)
Reviewer for CogSci (2020, 2021, 2022)
Reviewer for ICML 2024 Workshop LLMs and Cognition
Action Editor for ACL Rolling Review (2023, 2024)
Student representative at the MIT BCS Seminars Committee (2020–2022)
Mentor in Harvard PPREP Program

Awards

Harvard Psychology Hodgson Innovation Fund	2023

Honors

Singleton Fellowship	2019
Singleton Fellowship	2017
National Scholarship	2016
Star Graduate of Fudan University	2015
Honours Graduate Student of Shanghai	2015
First Prize of the 28^{th} Guanghua Self-independence Award [1]	2015
National Scholarship	2013
Hong Kong-Love Scholarship for Outstanding Student	2012
Freshman Scholarship of Fudan University	2011

[1] Guanghua Self-Independence Award is established by Fudan student company, and given to the students who have influential and leading achievement in one or more fields with perseverance, self-independence, and potentials for leadership.

Hobbies and skills	Violin Chinese calligraphy Traditional Chinese painting