Kenton Murray, Ph.D.

kenton@jhu.edu @kentonmurray kentonmurray
https://maieutic-nlp.github.io/website
https://www.kentonmurray.com

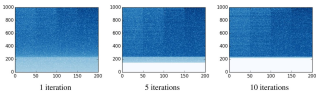

Education

2014 – 2020 **Ph.D., University of Notre Dame** Computer Science and Engineering.
Thesis title: *Learning Hyperparameters for Neural Machine Translation.*
Advisor: David Chiang
Committee: *Alexander M. Rush, Walter Scheirer, Meng Jiang*

2011 – 2013 **M.Sc. Computer Science, Carnegie Mellon University** Language Technologies
Thesis title: *A Semantic Scan Statistic for Novel Disease Outbreak Detection.*
Advisor: Daniel B. Niell

2005 – 2009 **BSE. Princeton University** Computer Science.
Thesis title: *Summarization by Latent Dirichlet Allocation: Superior Sentence Extraction through Topic Modeling.*
Advisor: David Blei

Professional Experience

2021 – Present **Research Scientist** Johns Hopkins University Human Language Technology Center of Excellence

Research Scientist Johns Hopkins University Center for Language and Speech Processing

Assistant Research Scientist Johns Hopkins University Department of Computer Science

2020 **Postdoctoral Fellow** Johns Hopkins University Center for Language and Speech Processing

2014 – 2019 **Graduate Research Assistant** University of Notre Dame NLP Lab

Summer 2016 **Research Intern** Allen Institute for Artificial Intelligence

2013-2014 **Research Associate** Qatar Computing Research Institute

2011-2013 **Research Assistant** Carnegie Mellon University Language Technology Institute

2009-2010 **Software Development Engineer in Test** Microsoft

Publications

2025

1. I. Abdulmumin, V. Agostinelli, T. Alumäe, *et al.*, "Findings of the IWSLT 2025 evaluation campaign," in *Proceedings of the 22nd International Conference on Spoken Language Translation (IWSLT 2025)*, E. Salesky, M. Federico, and A. Anastasopoulos, Eds., Vienna, Austria (in-person and online): Association for Computational Linguistics, Jul. 2025, pp. 412–481, ISBN: 979-8-89176-272-5. DOI: 10.18653/v1/2025.iwslt-1.44.

2. N. Bafna, E. Chang, N. R. Robinson, *et al.*, "Dialup! modeling the language continuum by adapting models to dialects and dialects to models," *arXiv preprint arXiv:2501.16581*, 2025.

3. N. Bafna, T. Li, K. Murray, *et al.*, "The translation barrier hypothesis: Multilingual generation with large language models suffers from implicit translation failure," *arXiv preprint arXiv:2506.22724*, 2025.

4 D. DeGenaro, E. Yang, D. Etter, *et al.*, "Fortify: Generative model fine-tuning with orpo for retrieval expansion of informal noisy text," in *Proceedings of the 1st Workshop on Multimodal Augmented Generation via Multimodal Retrieval (MAGMaR 2025)*, 2025, pp. 100–115.

5 E. Kennedy, S. Vadlamani, H. M. Lindsey, *et al.*, "Linking symptom inventories using semantic textual similarity," *Journal of Neurotrauma*, 2025.

6 T. Kocmi, E. Artemova, E. Avramidis, *et al.*, "Findings of the wmt25 general machine translation shared task: Time to stop evaluating on easy test sets," in *Proceedings of the Tenth Conference on Machine Translation, China. Association for Computational Linguistics*, 2025.

7 T. Kocmi, E. Avramidis, R. Bawden, *et al.*, "Preliminary ranking of wmt25 general machine translation systems," *arXiv preprint arXiv:2508.14909*, 2025.

8 R. Kriz and K. Murray, "Proceedings of the 1st workshop on multimodal augmented generation via multimodal retrieval (magmar 2025)," in *Proceedings of the 1st Workshop on Multimodal Augmented Generation via Multimodal Retrieval (MAGMaR 2025)*, 2025.

9 R. Kriz, K. Sanders, D. Etter, *et al.*, "Multivent 2.0: A massive multilingual benchmark for event-centric video retrieval," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 24 149–24 158.

10 T. Li, H. Xu, W. Tan, K. Murray, and D. Khashabi, "Upsample or upweight? balanced training on heavily imbalanced datasets," in *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2025, pp. 3325–3343.

11 P. McNamee, K. Duh, C. Carpenter, R. Colaianni, N. King, and K. Murray, "Oj4ocrmt: A large multilingual dataset for ocr-mt evaluation," in *Proceedings of Machine Translation Summit XX: Volume 1*, 2025, pp. 113–125.

12 A. Reddy, A. Martin, E. Yang, *et al.*, "Video-colbert: Contextualized late interaction for text-to-video retrieval," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 19 691–19 701.

13 N. R. Robinson, C. B. Monroc, R. Dent, *et al.*, "Findings of the first shared task for creole language machine translation at wmt25," in *Proceedings of the Tenth Conference on Machine Translation*, 2025.

14 N. R. Robinson, N. Bafna, X. He, *et al.*, "Jhu iwslt 2025 low-resource system description," in *Proceedings of the 22nd International Conference on Spoken Language Translation (IWSLT 2025)*, 2025, pp. 315–323.

15 S. Samuel, D. DeGenaro, J. Guallar-Blasco, *et al.*, "Mmmorrf: Multimodal multilingual modularized reciprocal rank fusion," in *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2025, pp. 4004–4009.

16 N. Sharma, K. Murray, and Z. Xiao, "Faux polyglot: A study on information disparity in multilingual large language models," in *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2025, pp. 8090–8107.

17 N. Verma, K. Murray, and K. Duh, "Dotresize: Reducing llm width via discrete optimal transport-based neuron merging," *arXiv preprint arXiv:2507.04517*, 2025.

18 N. Verma, K. Murray, and K. Duh, "Merging feed-forward sublayers for compressed transformers," *arXiv preprint arXiv:2501.06126*, 2025.

19 C. Xiao, M. Wiesner, D. Chakraborty, *et al.*, "Whisper-ut: A unified translation framework for speech and text," *arXiv preprint arXiv:2509.16375*, 2025.

20 H. Xu, K. Murray, P. Koehn, H. Hoang, A. Eriguchi, and H. Khayrallah, "X-ALMA: plug & play modules and adaptive rejection for quality translation at scale," in *The Thirteenth International Conference on*

Learning Representations, ICLR 2025, Singapore, April 24-28, 2025, OpenReview.net, 2025. URL: https://openreview.net/forum?id=csbf1p8xUq.

2024

1. I. S. Ahmad, A. Anastasopoulos, O. Bojar, *et al.*, "Findings of the iwslt 2024 evaluation campaign," *arXiv preprint arXiv:2411.05088*, 2024.
2. N. Bafna, K. Murray, and D. Yarowsky, "Evaluating large language models along dimensions of language variation: A systematik invesdigatiom uv cross-lingual generalization," pp. 18 742–18 762, 2024.
3. P. Foley, M. Wiesner, B. Odoom, L. P. G. Perera, K. Murray, and P. Koehn, "Where are you from? geolocating speech and applications to language identification," in *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2024, pp. 5114–5126.
4. T. Kocmi, E. Avramidis, R. Bawden, *et al.*, "Findings of the wmt24 general machine translation shared task: The llm era is here but mt is not solved yet," in *Proceedings of the Ninth Conference on Machine Translation*, 2024, pp. 1–46.
5. T. Kocmi, E. Avramidis, R. Bawden, *et al.*, "Preliminary wmt24 ranking of general mt systems and llms," *arXiv preprint arXiv:2407.19884*, 2024.
6. T. Li, H. Xu, P. Koehn, D. Khashabi, and K. Murray, "Error norm truncation: Robust training in the presence of data noise for text generation models," in *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*, OpenReview.net, 2024. URL: https://openreview.net/forum?id=zMvMwNvs4R.
7. B. B. Odoom, L. P. G. Perera, P. Hansanti, *et al.*, "Speech data from radio broadcasts for low resource languages," in *Proceedings of the 21st International Conference on Spoken Language Translation (IWSLT 2024)*, 2024, pp. 134–139.
8. B. B. Odoom, N. Robinson, E. Rippeth, *et al.*, "Can synthetic speech improve end-to-end conversational speech translation?" In *Proceedings of the 16th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, 2024, pp. 167–177.
9. N. Robinson, R. Dabre, A. Shurtz, *et al.*, "Kreyòl-mt: Building mt for latin american, caribbean and colonial african creole languages," in *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2024, pp. 3083–3110.
10. N. R. Robinson, K. Sun, C. Xiao, *et al.*, "Jhu iwslt 2024 dialectal and low-resource system description," in *Proceedings of the 21st International Conference on Spoken Language Translation (IWSLT 2024)*, 2024, pp. 140–153.
11. I. Said Ahmad, A. Anastasopoulos, O. Bojar, *et al.*, "Findings of the iwslt 2024 evaluation campaign," *arXiv e-prints*, arXiv–2411, 2024.
12. W. Tan, H. Xu, L. Shen, *et al.*, "Narrowing the gap between zero-and few-shot machine translation by matching styles," in *Findings of the Association for Computational Linguistics: NAACL 2024*, 2024, pp. 490–502.
13. N. Verma, K. Murray, and K. Duh, "Exploring geometric representational disparities between multilingual and bilingual translation models," in *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, Torino, Italy: ELRA and ICCL, May 2024, pp. 6909–6921. URL: https://aclanthology.org/2024.main-1.604.
14. H. Xu, A. Sharaf, Y. Chen, *et al.*, "Contrastive preference optimization: Pushing the boundaries of LLM performance in machine translation," in *Forty-first International Conference on Machine Learning, ICML*

2024, Vienna, Austria, July 21-27, 2024, OpenReview.net, 2024. URL: https://openreview.net/forum?id=51iwkioZpn.

2023

1. M. Agarwal, S. Agarwal, A. Anastasopoulos, *et al.*, "Findings of the iwslt 2023 evaluation campaign," Association for Computational Linguistics, 2023.
2. S. Barham, O. Weller, M. Yuan, *et al.*, "Megawika: Millions of reports and their sources across 50 diverse languages," *arXiv preprint arXiv:2307.07049*, 2023.
3. K. Dhole, V. Gangal, S. Gehrmann, *et al.*, "Nl-augmenter: A framework for task-sensitive natural language augmentation," in *Northern European Journal of Language Technology, Volume 9*, 2023.
4. T. Kocmi, E. Avramidis, R. Bawden, *et al.*, "Findings of the 2023 conference on machine translation (wmt23): Llms are here but not quite there yet," in *WMT23-Eighth Conference on Machine Translation*, 2023, pp. 198–216.
5. T. Li and K. Murray, "Why does zero-shot cross-lingual generation fail? an explanation and a solution," in *Findings of the Association for Computational Linguistics: ACL 2023*, 2023, pp. 12 461–12 476.
6. H. Xu, M. Elbayad, K. Murray, J. Maillard, and V. Goswami, "Towards being parameter-efficient: A stratified sparsely activated transformer with dynamic capacity," in *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023, pp. 12 858–12 870.
7. H. Xu, W. Tan, S. Li, *et al.*, "Condensing multilingual knowledge with lightweight language-specific modules," in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2023, pp. 1575–1587.

2022

1. A. Anastasopoulos, L. Barrault, L. Bentivogli, *et al.*, "Findings of the iwslt 2022 evaluation campaign.," in *Proceedings of the 19th international conference on spoken language translation (IWSLT 2022)*, Association for Computational Linguistics, 2022, pp. 98–157.
2. S. S. Li and K. Murray, "Language agnostic code-mixing data augmentation by predicting linguistic patterns," *arXiv preprint arXiv:2211.07628*, 2022.
3. S. Nair, E. Yang, D. Lawrie, *et al.*, "Transfer learning approaches for building cross-language dense retrieval models," in *European Conference on Information Retrieval*, Springer International Publishing Cham, 2022, pp. 382–396.
4. N. Verma, K. Murray, and K. Duh, "Strategies for adapting multilingual pre-training for domain-specific machine translation," in *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, 2022, pp. 31–44.
5. H. Xu, P. Koehn, and K. Murray, "The importance of being parameters: An intra-distillation method for serious gains," in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 2022, pp. 170–183.
6. H. Xu and K. Murray, "Por qué não utiliser alla språk? mixed training with gradient optimization in few-shot cross-lingual transfer," in *Findings of the Association for Computational Linguistics: NAACL 2022*, 2022, pp. 2043–2059.

2021

1. T. Q. Nguyen, K. Murray, and D. Chiang, "Data augmentation by concatenation for low-resource translation: A mystery and a solution," in *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, 2021, pp. 287–293.

2. E. Stengel-Eskin, K. Murray, S. Zhang, A. S. White, and B. Van Durme, "Joint universal syntactic and semantic parsing," *Transactions of the Association for Computational Linguistics*, vol. 9, pp. 756–773, 2021.
3. H. Xu, S. Ebner, M. Yarmohammadi, A. S. White, B. Van Durme, and K. Murray, "Gradual fine-tuning for low-resource domain adaptation," in *Proceedings of the Second Workshop on Domain Adaptation for NLP*, 2021, pp. 214–221.
4. H. Xu, B. Van Durme, and K. Murray, "Bert, mbert, or bibert? a study on contextualized embeddings for neural machine translation," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 6663–6675.
5. M. Yarmohammadi, S. Wu, M. Marone, *et al.*, "Everything is all it takes: A multipronged strategy for zero-shot cross-lingual information extraction," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 1950–1967.

2020

1. H. Khayrallah, J. Bremerman, A. D. McCarthy, K. Murray, W. Wu, and M. Post, "The jhu submission to the 2020 duolingo shared task on simultaneous translation and paraphrase for language education," in *Proceedings of the Fourth Workshop on Neural Generation and Translation*, 2020.
2. S. Mayhew, K. Bicknell, C. Brust, B. McDowell, W. Monroe, and B. Settles, "Simultaneous translation and paraphrase for language education," in *Proceedings of the ACL Workshop on Neural Generation and Translation (WNGT). ACL*, 2020.
3. A. Poliak, M. Fleming, C. Costello, *et al.*, "Collecting verified covid-19 question answer pairs," in *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, 2020.

2019

1. K. Murray, "Learning hyperparameters for neural machine translation," Ph.D. dissertation, University of Notre Dame, 2019.
2. K. Murray, B. DuSell, and D. Chiang, "Efficiency through auto-sizing: Notre dame nlp's submission to the wngt 2019 efficiency task," in *Proceedings of the 3rd Workshop on Neural Generation and Translation*, arXiv preprint arXiv:1910.07134, 2019, pp. 297–301.
3. K. Murray, J. Kinnison, T. Q. Nguyen, W. Scheirer, and D. Chiang, "Auto-sizing the transformer network: Improving speed, efficiency, and performance for low-resource machine translation," in *Proceedings of the 3rd Workshop on Neural Generation and Translation*, arXiv preprint arXiv:1910.06717, 2019, pp. 231–240.
4. K. W. Murray and J. Krishnamurthy, *System and method for program induction using probabilistic neural programs*, US Patent App. 16/044,220, Jan. 2019.

2018

1. U. Hermjakob, Q. Li, D. Marcu, *et al.*, "Incident-driven machine translation and name tagging for low-resource languages," *Machine Translation*, vol. 32, pp. 59–89, 2018.
2. K. Murray and D. Chiang, "Correcting length bias in neural machine translation," in *Proceedings of the Third Conference on Machine Translation*, 2018, pp. 212–223.
3. X. Zhang, G. Kumar, H. Khayrallah, *et al.*, "An empirical exploration of curriculum learning for neural machine translation," *arXiv preprint arXiv:1811.00739*, 2018.

2017

1. L. Cheung, T. Gowda, U. Hermjakob, *et al.*, "Elisa system description for lorehlt 2017," *Proc. Low Resource Human Lang. Technol*, pp. 51–59, 2017.

2016

1. D. Chiang and K. Murray, "Auto-sizing neural networks: With applications to n-gram language models," 2016.
2. A. Maurya, K. Murray, Y. Liu, C. Dyer, W. W. Cohen, and D. B. Neill, "Semantic scan: Detecting subtle, spatially localized events in text streams," *arXiv preprint arXiv:1602.04393*, 2016.
3. K. W. Murray and J. Krishnamurthy, "Probabilistic neural programs," *arXiv preprint arXiv:1612.00712, Neural Abstract Machines and Program Induction (NAMPI)*, 2016.

2015

1. K. Murray and D. Chiang, "Auto-sizing neural networks: With applications to n-gram language models," in *Empirical Methods in Natural Language Processing (EMNLP)*, 2015.

2014

1. A. Beutel, K. Murray, C. Faloutsos, and A. J. Smola, "Cobafi — collaborative bayesian filtering," in *Proceedings of the 23rd International World Wide Web Conference*, 2014.

2013

1. W. Ammar, V. Chahuneau, M. Denkowski, *et al.*, "The cmu machine translation systems at wmt 2013: Syntax, synthetic translation options, and pseudo-references," in *Proceedings of the Eighth Workshop on Statistical Machine Translation*, 2013, pp. 70–77.
2. K. Murray, "A semantic scan statistic for novel disease outbreak detection," Ph.D. dissertation, Carnegie Mellon University, 2013.
3. H. Sajjad, F. Guzmán, P. Nakov, *et al.*, "Qcri at iwslt 2013: Experiments in arabic-english and english-arabic spoken language translation," 2013.

2012

1. K. W. Murray and N. Orii, "Automatic essay scoring," *IEICE Transactions on Information and Systems*, vol. 102, no. 1, pp. 147–155, 2012.

2009

1. K. W. Murray, "Summarization by latent dirichlet allocation: Superior sentence extraction through topic modeling," *A senior thesis for Bachelors degree, Princeton University*, 2009.

Graduate Student Advising

2021-2024	**Haoran Xu (PhD)** Co-Advised with Philipp Koehn
2021-Present	**Neha Verma (PhD Anticipated)** Co-Advised with Kevin Duh
2023-Present	**Nathaniel R. Robinson (PhD Anticipated)** Co-Advised with Sanjeev Khudanpur
2023-2024	**Tianjian Li (MS)**
	Ujvala Pradeep (MS)
2022-2023	**Stella Li (MS)**
2024-Present	**Dengjia Zhang (MS Anticipated)**
	Aidan Alme (MS Anticipated)

Graduate Student Advising (continued)

2025-Present — **Xiaojian Sun (MS Anticipated)**

Teaching

Fall 2025 — **Explaining AI through Games** *Johns Hopkins University* EN.501.145

Data Science Meets Political Science *Johns Hopkins University* AS.190.360

Spring 2023 — **Advanced NLP: Multilingual Methods** *Johns Hopkins University* EN.601.764

Multilingual Methods in NLP *Georgetown University* LING4466

Spring 2022 — **Advanced NLP: Multilingual Methods** *Johns Hopkins University* EN.601.764

Spring 2017 — **Westville Education Initiative** *Westville Prison (Volunteer TA)*

Professional Activities

2025-Present — **President, ACL SIGIL** Founder and first president of the ACL Special Interest Group on Image and Language (SIGIL)

2026 — **ACL Tutorial Co-Chair** Select and organize tutorials for 2026 Association of Computational Linguistics Conference

2025 — **NAACL Workshops Co-Chair** Select and organize workshops for the 2025 Nations of the Americas Association of Computational Linguistics Conference

Co-Organizer WMDQS

2025, 2026 — **MAGMaR Co-Editor** Co-founder and organizer of the 1st and 2nd Workshop on Multimodal Retrieval Augmented Generation

2024 — **Senior Area Chair, NAACL**

Co-Lead SCALE 2024 10 week, 40+ person summer workshop on video event retrieval

2023 — **Co-Lead SCALE 2023** 10 week, 40+ person summer workshop on speech translation

2020, 2021 — **AMTA Student Mentorship** Organized mentoring opportunities for students at the conference

2018 — **Midwest Speech and Language Days Local Organizer** Non-archival conference.

2016 — **MTMA Local Organizing Committee** Week-long hackathon focused on improving machine translation

2024-Present — **Faculty Co-Mentor MASC-ALL** Annual, non-archival conference geared towards PhD students to present work to their peers

Invited Talks

2024 — **Video-based Event Retrieval** *ETH Zurich*

Improving Preference Optimization and Pre-Training Algorithms: Insights from focusing on Multilingual NLP and Machine Translation *University of Washington (Tsvetshop Lab)*

2021 — **Learning Hyperparameters for Neural Machine Translation** *Clemson University*

2019 — **Learning Neural Network Hyperparameters for Machine Translation** *University of Southern California Information Sciences Institute*

Learning Neural Network Hyperparameters for Machine Translation *Army Research Lab*

Learning Neural Network Hyperparameters for Machine Translation *Johns Hopkins University*

Invited Talks (continued)

2017 **Learning Neural Network Structure for Natural Language** *University of Southern California Information Sciences Institute*

Miscellaneous Experience

PhD Committees

2024 **Xuan Zhang** Advisor: *Kevin Duh*

2022 **Patrick Xia** Advisor: *Benjamin Van Durme*

Graduate Board Oral Exams

2025 **Weiting (Steven) Tan** Advisor: *Philipp Koehn*

2023 **Xuan Zhang** Advisor: *Kevin Duh*

2022 **Seth Ebner** Advisor: *Benjamin Van Durme*

YUnmo Chen Advisor: *Benjamin Van Durme*

Local Service

2021-Present **CLSP PhD Admissions Chair**

Miscellaneous

2008-2009 **Princeton Charter Eating Club Officer**

2005-2007 **Princeton Varsity Crew Team**