Jungo Kasai

Contact Information

Kotoba Technologies, Inc. jkasai@kotoba.tech
Toyota Technological Institute at Chicago jkasai@ttic.edu
6045 S. Kenwood Ave.
Chicago, IL 60637, USA https://homes.cs.washington.edu/~jkasai/

Employment Experience

Co-founder and Chief Technology Officer 07/2023 - *present*
Kotoba Technologies Inc.

Research Assistant Professor 09/2023 - *present*
Toyota Technological Institute at Chicago (TTIC)

Research Assistant with Noah A. Smith 09/2018 - 08/2023
University of Washington

Research Intern with Keisuke Sakaguchi and Ronan Le Bras 07/2021 - 09/2022
Allen Institute for Artificial Intelligence

Research Intern with Yizhe Zhang and Yi Mao 07/2020 - 09/2020
Microsoft AI

Research Intern with James Cross 09/2019 - 01/2020
Facebook AI

Research Intern with Yunyao Li and Kun Qian 05/2018 - 09/2018
IBM Research-Almaden

Research Assistant with Dragomir Radev, Robert Frank, and Owen Rambow 05/2016 - 05/2018
Yale University

Education

University of Washington, Seattle, WA 2018 - 2023
Ph.D. in Computer Science and Engineering
Advisor: Noah A. Smith

Yale University, New Haven, CT 2013 - 2017
B.S. in Statistics
Magna Cum Laude, *Phi Beta Kappa*, *Distinction in the Major*

Honors and Awards

Best Paper Award at NAACL, 2022

IBM Ph.D. Fellowship, 2020

Masason Foundation Scholarship, 2018

Funai Overseas Scholarship, 2017

Richter Fellowship, Yale University, 2016

Herb Scarf Summer Research Fellowship, Yale University, 2016

Richard U. Light Fellowship, Yale University, 2014-2015

Gold Medal in the Japanese Philosophy Olympiad, 2011-2012

Invited Talks

Dramatic Five Years of AI and NLP: The Present and Future of Large Language Models
– University of Tokyo. October, 2023.
– Nara Institute of Science and Technology. October, 2023.

Democratized Natural Language Processing: Efficiency, Evaluation, and Beyond
– Apple Knowledge Platform. December, 2022.
– Tohoku University. January, 2022.
– Toyota Technological Institute at Chicago. February, 2023.
– Microsoft Research. February, 2023.
– Google DeepMind. March, 2023.
– TikTok USA. April, 2023.
– Stony Brook University. April, 2023.

Evaluating GPT-4 and ChatGPT on Japanese Medical Licensing Examinations
– Quantum Photonics. May, 2023.

Transparent and Dynamic Evaluation for Natural Language Processing
– Apple Intelligent System Experience. September, 2022.

Twist Decoding: Diverse Generators Guide Each Other
– University of Lisbon/Unbabel. June, 2022.

Bidimensional Leaderboards: Generate and Evaluate Language Hand in Hand
– Georgia Institute of Technology. December, 2021.

Finetuning Pretrained Transformers into RNNs
– Microsoft. April, 2021.
– MLOps, Production & Engineering New York. April, 2021.

Deep Encoder, Shallow Decoder: Reevaluating Non-autoregressive Machine Translation
– NLP Colloquium, Japan. July, 2021.

Low-resource Deep Entity Resolution
– IBM AI Horizons Seminar Series. September, 2019.

Neural Network-based Supertagging and Parsing of Tree-Adjoining Grammars
– Heinrich Heine University of Düsseldorf. December, 2018.
– Interactions Lab. December, 2017.

Mentoring

Haoxin (Alan) Li [3] – B.S., UW CSE 05/2022 - *present*
Hongjin Su [5] – B.S., the Chinese University of Hong Kong 10/2021 - *present*
→ Ph.D. student at the University of Hong Kong CS
Chi-Liang Liu – M.S., National Taiwan University 09/2022 - 04/2023
Jiayi Xin [5] – B.S., University of Hong Kong CS 10/2021 - 04/2022
Lavinia Dunagan [9, 10] – B.S., UW CSE 09/2021 - 06/2022
→ Ph.D. student at the University of Michigan iSchool
Jacob Morrison [9, 10] – M.S., UW Linguistics 09/2021 - 06/2022
→ Predoctoral young investigator at the Allen Institute for AI
Ivan Montero [8] – B.S., UW CSE 06/2021 - 06/2022
→ Machine learning engineer at Apple
Ryan Li – B.S., UW CSE 03/2022 - 06/2022
Xinyan Velocity Yu [15, 17, 33] – B.S., UW CSE 09/2020 - 04/2021
→ AI resident at Meta AI → Ph.D. student at University of Southern California CS
Leo (Zeyu) Liu [14] – M.S., UW CSE 02/2020 - 06/2021
→ AI resident at Meta AI → Ph.D. student at UT Austin CS
Chi-Hsun (Sean) Hung – B.S., UW CSE 10/2018 - 05/2019

Teaching Experience

Teaching Assistant at the Paul G. Allen School of CSE, UW
NLP Capstone Spring 2021
Natural Language Processing (Graduate Level) Winter 2019

Guest Lecturer at the Simpson Center for the Humanities, UW
Machine Translation: Its History, Revolution, and Future Prospects Spring 2021

Guest Lecturer at the Paul G. Allen School of CSE, UW
Dependency Parsing Winter 2019

Teaching Assistant at the Department of CS, Yale University
Natural Language Processing Spring 2017

Professional Service

Program comittee member/reviewer: ACL Rolling Review (2021–*present*), ACL (2021–2022), EMNLP (2021–2022), NAACL (2022), NeurIPS (2021–2022), ICML (2021–2022), ICLR (2022–2023), MIA (2022).

Refereed Conference Publications

[1] Yushi Hu, Benlin Liu, **Jungo Kasai**, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A. Smith. TIFA: Accurate and Interpretable Text-to-Image Faithfulness Evaluation with Question Answering In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023.

[2] Zheng-Xin Yong, Hailey Schoelkopf, Niklas Muennighoff, Alham Fikri Aji, David Ifeoluwa Adelani, Khalid Almubarak, M Saiful Bari, Lintang Sutawika, **Jungo Kasai**, Ahmed Baruwa, Genta Indra Winata, Stella Biderman, Edward Raff, Dragomir Radev, and Vassilina Nikoulina. BLOOM+1: Adding Language Support to BLOOM for Zero-Shot Prompting . In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.

[3] Haoxin Li, Phillip Keung, Daniel Cheng, **Jungo Kasai**, and Noah A. Smith. NarrowBERT: Accelerating Masked Language Model Pretraining and Inference. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.

[4] Hongjin Su, Weijia Shi, **Jungo Kasai**, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A Smith, Luke Zettlemoyer, and Tao Yu. One Embedder, Any Task: Instruction-Finetuned Text Embeddings . In *Findings of the Association for Computational Linguistics: ACL 2023*, 2023.

[5] Hongjin Su, **Jungo Kasai**, Chen Henry Wu, Weijia Shi, Tianlu Wang, Jiayi Xin, Rui Zhang, Mari Ostendorf, Luke Zettlemoyer, Noah A. Smith, and Tao Yu. Selective Annotation Makes Language Models Better Few-Shot Learners. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

[6] **Jungo Kasai**, Keisuke Sakaguchi, Ronan Le Bras, Hao Peng, Ximing Lu, Dragomir Radev, Yejin Choi, and Noah A. Smith. Twist Decoding: Diverse Generators Guide Each Other. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2022.

[7] Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, **Jungo Kasai**, Yejin Choi, Noah A. Smith, and Daniel S. Weld. GENIE: Toward Reproducible and Standardized Human Evaluation for Text Generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2022.

[8] Michael Hassid, Hao Peng, Daniel Rotem, **Jungo Kasai**, Ivan Montero, Noah A. Smith, and Roy Schwartz. How Much Does Attention Actually Attend? Questioning the Importance of Attention in Pretrained Transformers. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, 2022.

[9] **Jungo Kasai**, Keisuke Sakaguchi, Ronan Le Bras, Lavinia Dunagan, Jacob Morrison, Alexander R. Fabbri, Yejin Choi, and Noah A. Smith. Bidimensional Leaderboards: Generate and Evaluate Language Hand in Hand. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistic (NAACL)*, 2022.

[10] **Jungo Kasai**, Keisuke Sakaguchi, Lavinia Dunagan, Jacob Morrison, Ronan Le Bras, Yejin Choi, and Noah A. Smith. Transparent Human Evaluation for Image Captioning. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistic (NAACL)*, 2022.

[11] Ximing Lu, Sean Welleck, Peter West, Liwei Jiang, **Jungo Kasai**, Daniel Khashabi, Ronan Le Bras, Lianhui Qin, Youngjae Yu, Rowan Zellers, Noah A. Smith, and Yejin Choi. NeuroLogic A*esque Decoding: Constrained Text Generation with Lookahead Heuristics. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistic (NAACL)*, 2022. **Best Paper Award.**

[12] Hao Peng, **Jungo Kasai**, Nikolaos Pappas, Dani Yogatama, Zhaofeng Wu, Lingpeng Kong, Roy Schwartz, and Noah A. Smith. ABC: Attention with Bounded-memory Control. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022.

[13] **Jungo Kasai**, Hao Peng, Yizhe Zhang, Dani Yogatama, Gabriel Ilharco, Nikolaos Pappas, Yi Mao, Weizhu Chen, and Noah A. Smith. Finetuning Pretrained Transformers into RNNs. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021.

[14] Leo Z. Liu, Yizhong Wang, **Jungo Kasai**, Hannaneh Hajishirzi, and Noah A. Smith Probing Across Time: What Does RoBERTa Know and When?. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, 2021.

[15] Akari Asai, Xinyan Yu, **Jungo Kasai**, and Hanna Hajishirzi. One Question Answering Model for Many Languages with Cross-lingual Dense Passage Retrieval. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

[16] **Jungo Kasai**, Nikolaos Pappas, Hao Peng, James Cross, and Noah A. Smith. Deep Encoder, Shallow Decoder: Reevaluating the Speed-Quality Tradeoff in Machine Translation. In *Proceedings of the Conference of the International Conference on Learning Representations (ICLR)*, 2021.

[17] Akari Asai, **Jungo Kasai**, Jonathan H. Clark, Kenton Lee, Eunsol Choi, and Hannaneh Hajishirzi. XOR QA: Cross-lingual Open-Retrieval Question Answering. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistic (NAACL)*, 2021.

[18] Farhad Akhbardeh, Arkady Arkhangorodsky, Magdalena Biesialska, Ondřej Bojar, Rajen Chatterjee, Vishrav Chaudhary, Marta R. Costa-jussa, Cristina España-Bonet, Angela Fan, Christian Federmann, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Leonie Harter, Kenneth Heafield, Christopher Homan, Matthias Huck, Kwabena Amponsah-Kaakyire, **Jungo Kasai**, Daniel Khashabi, Kevin Knight, Tom Kocmi, Philipp Koehn, Nicholas Lourie, Christof Monz, Makoto Morishita, Masaaki Nagata, Ajay Nagesh, Toshiaki Nakazawa, Matteo Negri, Santanu Pal, Allahsera Auguste Tapo, Marco Turchi, Valentin Vydrin, and Marcos Zampieri. Findings of the Conference on Machine Translation (WMT21). In *Proceedings of the Conference on Machine Translation (WMT)*, 2021.

[19] **Jungo Kasai**, James Cross, Marjan Ghazvininejad, and Jiatao Gu. Non-Autoregressive Machine Translation with Disentangled Context Transformer. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.

[20] Phoebe Mulcaire,∗ **Jungo Kasai**,∗ and Noah A. Smith. Low-Resource Parsing with Crosslingual Contextualized Representations. In *Proceedings of the Conference on Computational Natural Language Learning (CoNLL)*, 2019. ∗ = equal contribution.

[21] **Jungo Kasai**, Kun Qian, Sairam Gurajada, Yunyao Li, and Lucian Popa. Low-resource Deep Entity Resolution with Transfer and Active Learning. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.

[22] **Jungo Kasai**, Dan Friedman, Robert Frank, Dragomir Radev, and Owen Rambow. Syntax-aware Neural Semantic Role Labeling with Supertags. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistic (NAACL)*, 2019.

[23] Phoebe Mulcaire, **Jungo Kasai**, and Noah A. Smith. Polyglot Contextual Representations Improve Crosslingual Transfer. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistic (NAACL)*, 2019.

[24] Michihiro Yasunaga, **Jungo Kasai**, Rui Zhang, Alexander R. Fabbri, Irene Li, Dan Friedman, and Dragomir R. Radev. ScisummNet: A Large Annotated Corpus and Content-Impact Models for Scientific Paper Summarization with Citation Networks. In *Proceedings of the Conference of the Conference on Artificial Intelligence (AAAI)*, 2019.

[25] **Jungo Kasai** and Robert Frank. Jabberwocky Parsing: Dependency Parsing with Lexical Noise. In *Proceedings of the Society for Computation in Linguistics (SCiL)*, 2019.

[26] **Jungo Kasai**, Robert Frank, Pauli Xu, William Merrill, and Owen Rambow. End-to-end Graph-based TAG Parsing with Neural Networks. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistic (NAACL)*, 2018.

[27] Michihiro Yasunaga, **Jungo Kasai**, and Dragomir Radev. Robust Multilingual Part-of-Speech Tagging via Adversarial Training. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistic (NAACL)*, 2018.

[28] **Jungo Kasai**, Robert Frank, R. Thomas McCoy, Owen Rambow, and Alexis Nasr. TAG Parsing with Neural Networks and Vector Representations of Supertags. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.

Preprints

[29] **Jungo Kasai**, Yuhei Kasai, Keisuke Sakaguchi, Yutaro Yamada, and Dragomir Radev. Evaluating GPT-4 and ChatGPT on Japanese Medical Licensing Examinations Under review, 2023

[30] John J. Nay, David Karamardian, Sarah B. Lawsky, Wenting Tao, Meghana Bhat, Raghav Jain, Aaron Travis Lee, Jonathan H. Choi, and **Jungo Kasai**. Large Language Models as Tax Attorneys: A Case Study in Legal Capabilities Emergence Under review, 2023

[31] Zhoujun Cheng, **Jungo Kasai**, and Tao Yu. Batch Prompting: Efficient Inference with Large Language Model APIs Under review, 2023

[32] Orevaoghene Ahia, Sachin Kumar, Hila Gonen, **Jungo Kasai**, David R. Mortensen, Noah A. Smith, and Yulia Tsvetkov. Do All Languages Cost the Same? Tokenization in the Era of Commercial Language Models Under review, 2023.

[33] **Jungo Kasai**, Keisuke Sakaguchi, Yoichi Takahashi, Ronan Le Bras, Akari Asai, Xinyan Yu, Dragomir Radev, Noah A. Smith, Yejin Choi, and Kentaro Inui. RealTime QA: What's the Answer Right Now? Under review, 2022.

[34] **Jungo Kasai**, Keisuke Sakaguchi, Ronan Le Bras, Dragomir Radev, Yejin Choi, and Noah A. Smith. Beam Decoding with Controlled Patience. Under review, 2022.

[35] Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Luke Benson, Lucy Sun, Ekaterina Zubova, Yujie Qiao, Matthew Burtell, David Peng, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, Ansong Ni, Linyong Nan, **Jungo Kasai**, Tao Yu, Rui Zhang, Shafiq Joty, Alexander R. Fabbri, Wojciech Kryscinski, Xi Victoria Lin, Caiming Xiong, and Dragomir Radev. FOLIO: Natural Language Reasoning with First-Order Logic. Under review, 2022.

Refereed Workshop Publications

[36] Akari Asai, Shayne Longpre, **Jungo Kasai**, Chia-Hsuan Lee, Rui Zhang, Junjie Hu, Ikuya Yamada, Jonathan H. Clark, and Eunsol Choi. MIA 2022 Shared Task: Evaluating Cross-lingual Open-Retrieval Question Answering for 16 Diverse Languages. In *Proceedings the Workshop on Multilingual Information Access (MIA)*, 2022.

[37] Jeff Da and **Jungo Kasai**. Cracking the Contextual Commonsense Code: Understanding Commonsense Reasoning Aptitude of Deep Contextual Representations. In *Proceedings the Workshop on Commonsense Inference in Natural Language Processing (COIN)*, 2019.

[38] Dan Friedman,∗ **Jungo Kasai**,∗ R. Thomas McCoy,∗ Robert Frank, Owen Rambow, and Forrest Davis. Linguistically Rich Vector Representations of Supertags for TAG Parsing. In *Proceedings of the Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+)*, 2017.

[39] Pauli Xu, Robert Frank, Jungo Kasai, and Owen Rambow TAG Parsing Evaluation using Textual Entailment. In *Proceedings of the Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+)*, 2017.

Other Publications

[40] Yuhei Kasai, Takayuki Kitai, Junji Morita, Takuya Okada, **Jungo Kasai**, and Tsutomu Fujita. Successful Catheter Ablation for Verapamil-Sensitive Idiopathic Left Ventricular Tachycardia Guided by Dual Post-QRS Wave P1 Potentials After Catheter-Induced Mechanical Block. *Heart Rhythm Case Reports*, 2023.

[41] Yuhei Kasai, **Jungo Kasai**, Takuya Haraguchi, Takayuki Kitai, Junji Morita, Takuya Okadac, Masanaga Tsujimoto, and Tsutomu Fujita. Lock and dock: Two-step transvenous retrieval of a fractured femoral sheath with a vascular snare via the right internal jugular vein. *Journal of Cardiology Cases*, 2023.

[42] Takayuki Kitai, Yuhei Kasai, **Jungo Kasai**, Ryo Otake, Ryo Horita, Hidemasa Shitan, Junji Morita, and Tsutomu Fujita Iatrogenic Aortic Regurgitation After Catheter Ablation of Aortic Right-left Interleaflet Triangle. *JACC Clinical Electrophysiology*, 2023.

[43] Yuhei Kasai, **Jungo Kasai**, Takayuki Kitai, Ryo Horita, Junji Morita, and Tsutomu Fujita. Silent steam pop detected by transesophageal echocardiography for premature ventricular contractions originating from the aortomitral continuity. *Journal of Echocardiography*, 2023.

[44] Yuhei Kasai, **Jungo Kasai**, Takayuki Kitai, Junji Morita, and Tsutomu Fujita. Iatrogenic left ventricular pseudoaneurysm after successful radiofrequency catheter ablation for premature ventricular contraction originating from the posterior papillary muscles. *Circulation Journal*, 2022.

[45] Takayuki Kitai, Yuhei Kasai, **Jungo Kasai**, Junji Morita, and Tsutomu Fujita Mediastinal Hematoma Associated With Radiofrequency Catheter Ablation of Atrial Fibrillation. *Circulation Journal*, 2022.

[46] Yuhei Kasai, **Jungo Kasai**, So Asano, Takahiko Nagase, Yukio Sekiguchi, and Junichi Nitta. Removing pericardial drainage tube for acute cardiac tamponade associated with catheter ablation of atrial fibrillation can trigger hemoperitoneum from severe liver bleeding. *Heart Rhythm Case Reports*, 2022.

[47] Yuhei Kasai, **Jungo Kasai**, Syuichi Sahashi, Sandeep Shakya, Hiroki Kuji, Naoki Hayakawa, Kotaro Miyaji, and Junji Kanda Revisiting where to apply preimplant mapping to improve P-wave sensing of insertable cardiac monitors. *Journal of Arrhythmia*, 2022.

[48] Yuhei Kasai, **Jungo Kasai**, Yukio Sekiguchi, Takahiko Nagase, and Junichi Nitta. Idiopathic premature ventricular contractions originating from the distal Purkinje fiber network of the right bundle branch.. *Journal of Arrhythmia*, 2022.

[49] Yuhei Kasai, **Jungo Kasai**, Yukio Sekiguchi, So Asano, Hiroshi Fukunaga, Takahiko Nagase, and Junichi Nitta. Apple Watch facilitates single-session catheter ablation of coexisting atrioventricular nodal reentrant tachycardia and atrioventricular reentrant tachycardia. *Clinical Case Reports*, 2021.

Skills

Software: Python, TensorFlow, PyTorch, Theano, Lua, Torch7, Linux, Matlab, R
Languages: Japanese (native), Mandarin Chinese (intermediate proficiency)

References

Noah A. Smith (Doctoral Advisor)
Amazon Professor of Machine Learning — UW Paul G. Allen School of CSE
Senior Research Manager — Allen Institute for AI
nasmith@cs.washington.edu

Yejin Choi
Brett Helsel Professor — UW Paul G. Allen School of CSE
Senior Research Manager — Allen Institute for AI
MacArthur Fellow
yejin@cs.washington.edu

Yunyao Li
Head of Machine Learning — Apple Knowledge Platform
yunyaoli@apple.com