

Karen Livescu

Professor

Toyota Technological Institute at Chicago
6045 S. Kenwood Ave., Chicago, IL 60637
<http://www.ttic.edu/livescu>

RESEARCH INTERESTS

Speech and language processing, as well as related problems in machine learning. Specific recent interests: spoken language models, efficiency in spoken and written language models, self-supervised representation learning, learning from multiple modalities, language-universal speech processing, automatic sign language understanding.

EDUCATION

- ◇ **Massachusetts Institute of Technology**
Ph.D. in Electrical Engineering and Computer Science, September 2005.
Dissertation: *Feature-Based Pronunciation Modeling for Automatic Speech Recognition*
Advisor: Dr. Jim Glass
M.S. in Electrical Engineering and Computer Science, 1999.
Thesis: *Analysis and Modeling of Non-Native Speech for Automatic Speech Recognition*
Advisor: Dr. Jim Glass
- ◇ **Technion–Israel Institute of Technology**, Haifa, Israel.
Visiting student in Electrical Engineering, 1996–97.
- ◇ **Princeton University**
A.B. in Physics with Honors and Certificate in Linguistics, 1996.
Thesis: *Analysis of Human and Parrot Phonation Using an Energy Operator and Energy Separation Algorithm*
Advisor: Prof. Ken Steiglitz

APPOINTMENTS

- ◇ **Professor**, Toyota Technological Institute at Chicago, 2021–present
- ◇ **Associate Professor, part-time**, Computer Science Department, U. Chicago, 2017–present
- ◇ **Affiliated Faculty**, Data Science Institute, U. Chicago
- ◇ **Visiting Scholar**, Stanford University Linguistics Department, 2023–2024
- ◇ **Special Faculty Researcher**, Language Technologies Institute, Carnegie Mellon University, 2023–2024
- ◇ **Associate Professor**, Toyota Technological Institute at Chicago, 2017–2021
- ◇ **Assistant Professor**, Toyota Technological Institute at Chicago, 2008–2017
- ◇ **Assistant Professor, part-time**, Computer Science Department, U. Chicago, 2009–17
- ◇ **Visiting Scientist**, Weizmann Institute of Science, Winter 2011–12, 2010–11, 2008–09
- ◇ **Research Assistant Professor**, Toyota Technological Institute at Chicago, 2008
- ◇ **Clare Boothe Luce Postdoctoral Lecturer**, Electrical Engineering and Computer Science department, Massachusetts Institute of Technology, 2005–07

HONORS/AWARDS

- ◇ **Fellow of the Institute of Electrical and Electronics Engineers (IEEE)**, elected 2024
- ◇ **Best Paper** (with W. Chen, W. Zhang, Y. Peng, X. Li, J. Tian, J. Shi, X. Chang, S. Maiti, and S. Watanabe), EMNLP 2024
- ◇ **ICML Test of Time Runner-Up** for Andrew *et al.*, “Deep Canonical Correlation Analysis”, 2023
- ◇ **Fellow of the International Speech Communication Association (ISCA)**, elected 2021
- ◇ **IEEE Signal Processing Society Distinguished Lecturer**, 2021–22
- ◇ **Amazon AWS Machine Learning Research Award**, 2020
- ◇ **Best Student Paper** (with C.-M. Chien, M. Zhang, J.-C. Chou), ASRU 2023
- ◇ **Best Paper nominee** (with H. Shi, J. Mao, K. Gimpel), ACL 2019
- ◇ **Best Paper nominee** (with H. Kamper, G. Goldwater), ASRU 2017
- ◇ **Best Paper** (with L. Tu, K. Gimpel), ACL Workshop on Representation Learning for NLP (RepL4NLP) 2017
- ◇ **Best Paper** (with P. Swaroop Madhyastha, M. Bansal, K. Gimpel), ACL Workshop on Representation Learning for NLP (RepL4NLP) 2016
- ◇ **Best Student Paper nominee** (with T. Kim, W. Wang, H. Tang), ICASSP 2016
- ◇ **Best Paper nominee** (with H. Tang, W. Wang, K. Gimpel), ASRU 2015
- ◇ **Google Research Award**, 2015
- ◇ **Google Research Award**, 2014 (with M. Bansal, K. Gimpel)
- ◇ **Best Student Paper 2nd Place** (with K. Levin, K. Henry, A. Jansen), ASRU 2013
- ◇ **Best Student Paper** (with P. Jyothi, E. Fosler-Lussier), Interspeech 2012
- ◇ **Feinberg Visiting Faculty Fellowship**, Weizmann Institute, winter 2010–11
- ◇ **Clare Boothe Luce Post-Doctoral Fellowship**, 2005–07
- ◇ **NSF Graduate Research Fellowship**, 1997–2000

PUBLICATIONS

JOURNALS

1. A. Pasad, C.-M. Chien, S. Settle, and K. Livescu,
“What do self-supervised speech models know about words?,”
Trans. ACL 12:372–91, April 2024.
2. A. Srivastava *et al.*,
“Beyond the imitation game: Quantifying and extrapolating the capabilities of language models,”
Trans. Machine Learning Research 5:1–95, 2023.
3. A. Mohamed, H.-y. Lee, L. Borgholt, J. D. Havtorn, J. Edin, C. Igel, K. Kirchhoff, S.-W. Li, K. Livescu, L. Maaløe, T. N. Sainath, S. Watanabe,
“Self-supervised speech representation learning: A review,” *IEEE Journal of Selected Topics in Signal Processing* **16**(6):1179–1210, September 2022.

4. H. Kamper, G. Shakhnarovich, and K. Livescu,
“Semantic speech retrieval with a visually grounded model of untranscribed speech,”
IEEE/ACM Transactions on Audio, Speech, and Language Processing **27**(1):89–98, January 2019.
5. E. M. Mugler, M. C. Tate, K. Livescu, J. W. Templar, M. A. Goldrick, M. W. Slutzky,
“Differential representation of articulatory gestures and phonemes in motor, premotor,
and inferior frontal cortices,”
J. Neuroscience **38**(46):9803–9813, November 2018.
6. H. Tang, L. Lu, L. Kong, K. Gimpel, K. Livescu, C. Dyer, N. A Smith, and S. Renals,
“End-to-end neural segmental models for speech recognition,”
IEEE Transactions of Selected Topics in Signal Processing, December 2017.
7. T. Kim, J. Keane, W. Wang, H. Tang, G. Shakhnarovich, J. Riggle, D. Brentari, and K. Livescu,
“Lexicon-free fingerspelling recognition from video: Data, models, and signer adaptation,”
Computer Speech and Language **46**:209–232, November 2017.
8. K. Livescu, P. Jyothi, and E. Fosler-Lussier,
“Articulatory feature-based pronunciation modeling,”
Computer Speech and Language **36**:212–232, March 2016.
9. J. Wieting, M. Bansal, K. Gimpel, and K. Livescu,
“From paraphrase database to compositional paraphrase model and back”,
Trans. ACL **3**:345–358, June 2015.
10. K. Livescu, E. Fosler-Lussier, and F. Metze,
“Subword modeling for automatic speech recognition,”
Signal Processing Magazine **29**(6):44–57, November 2012.
11. K. Saenko, K. Livescu, J. Glass, and T. Darrell,
“Multistream articulatory feature-based models for visual speech recognition,”
IEEE Trans. Pattern Analysis and Machine Intelligence **31**(9):1700–1707, September 2009.
12. S. King, J. Frankel, K. Livescu, E. McDermott, K. Richmond, and M. Wester,
“Speech production knowledge in automatic speech recognition,”
Journal of the Acoustical Society of America **121**(2):723–742, February 2007.
13. T. J. Hazen, I. Lee Hetherington, H. Shu, and K. Livescu,
“Pronunciation modeling using a finite-state transducer representation,”
Speech Communication **46**(2):189–203, June 2005.

REFEREED CONFERENCE/WORKSHOP PROCEEDINGS

1. W. Chen, C. Meng, J. Shi, M. Bartelds, S.-H. Wang, H.-H. Wang, R. Mosquera, S. Hincapie, D. Jurafsky, A. Anastosopoulos, H.-y. Lee, K. Livescu, and S. Watanabe,
“The ML-SUPERB 2.0 Challenge: Towards Inclusive ASR Benchmarking for All Language Varieties,”
Interspeech 2025.
2. S. Gueuwou, X. Du, G. Shakhnarovich, K. Livescu, and A. H. Liu,
“SHuBERT: Self-Supervised Sign Language Representation Learning via Multi-Stream Cluster Prediction,”
ACL 2025.

3. S. Gueuwou, X. Du, G. Shakhnarovich, and K. Livescu,
 “SignMusketeers: An Efficient Multi-Stream Approach for Sign Language Translation at Scale,”
ACL 2025 Findings.
4. C.-y. Huang *et al.*,
 “Dynamic-SUPERB Phase-2: A collaboratively expanding benchmark for measuring the capabilities of spoken language models with 180 tasks,”
ICLR 2025.
5. Y. Li, K. Livescu, and J. Zhou,
 “Chunk-Distilled Language Modeling,”
ICLR 2025.
6. J. Rosas-Smith, M. Bartelds, R. Huang, L. P. García-Perera, K. Livescu, D. Jurafsky, and A. Field,
 “Constructing Datasets From Public Police Body Camera Footage,”
ICASSP 2025.
7. T. Srivastava, J.-C. Chou, P. Shroff, K. Livescu, and C. Graziul,
 “Speech Recognition for Analysis of Police Radio Communication,”
SLT 2024.
8. W. Chen, W. Zhang, Y. Peng, X. Li, J. Tian, J. Shi, X. Chang, S. Maiti, K. Livescu, and S. Watanabe,
 “Towards robust speech representation learning for thousands of languages,”
EMNLP 2024.
(Best Paper Award)
9. K. Choi, A. Pasad, T. Nakamura, S. Fukayama, K. Livescu, and S. Watanabe,
 “Self-supervised speech representations are more phonetic than semantic,”
Interspeech 2024.
10. K. Kim, Y.-T. Hsu, P. Sridhar, K. Livescu, and S. Watanabe,
 “Convolution-augmented parameter-efficient fine-tuning for speech recognition,”
Interspeech 2024.
11. J. Shi, S.-H. Wang, W. Chen, M. Bartelds, V. B. Kumar, J. Tian, X. Chang, D. Jurafsky, K. Livescu, H.-y. Lee, and S. Watanabe,
 “ML-SUPERB 2.0: Benchmarking multilingual speech models across modeling constraints, languages, and datasets,”
Interspeech 2024.
12. S. Shon, K. Kim, Y.-T. Hsu, P. Sridhar, S. Watanabe, and K. Livescu,
 “DiscreteSLU: A large language model with self-supervised discrete speech units for spoken language understanding,”
Interspeech 2024.
13. J. Tian, Y. Peng, W. Chen, K. Choi, K. Livescu, and S. Watanabe,
 “On the effects of heterogeneous data sources on speech-to-text foundation models,”
Interspeech 2024.
14. S. Arora, A. Pasad, C.-M. Chien, J. Han, R. Sharma, J.-w. Jung, H. Dharmyal, W. Chen, S. Shon, H.-y. Lee, K. Livescu, and S. Watanabe,
 “On the Evaluation of Speech Foundation Models for Spoken Language Understanding,”
ACL 2024 Findings.

15. F. Shi, K. Gimpel, and K. Livescu,
 “Structured Tree Alignment for Evaluation of (Speech) Constituency Parsing,”
ACL 2024.
16. D. Yunis, K. Kshitij Patel, S. Wheeler, P. H. Pamplona Savarese, G. Vardi, K. Livescu, M. Maire, and M. Walter,
 “Grokking, Rank Minimization and Generalization in Deep Learning,”
ICML 2024 Workshop on Mechanistic Interpretability.
17. S. Arora, H. Futami, J.-w. Jung, Y. Peng, R. Sharma, Y. Kashiwagi, E. Tsunoo, K. Livescu, and S. Watanabe,
 “UniverSLU: Universal Spoken Language Understanding for Diverse Tasks with Natural Language Instructions,”
NAACL 2024.
18. J.-C. Chou, C.-M. Chien, and K. Livescu,
 “AV2Wav: Diffusion-based re-synthesis from continuous self-supervised features for audio-visual speech enhancement,”
ICASSP 2024.
19. S. Shon, K. Kim, P. Sridhar, Y.-T. Hsu, S. Watanabe, and K. Livescu,
 “Generative Context-Aware Fine-Tuning of Self-Supervised Speech Models,”
ICASSP 2024.
20. J.-C. Chou, C.-M. Chien, W.-N. Hsu, K. Livescu, A. Babu, A. Conneau, A. Baevski, and M. Auli,
 “Toward joint language modeling for speech units and text,”
Findings of EMNLP 2023.
21. M. Sandoval-Castañeda, Y. Li, B. Shi, D. Brentari, K. Livescu, and G. Shakhnarovich,
 “TTIC’s Submission to WMT-SLT 23,”
Eighth Conference on Machine Translation (WMT) 2023.
(Top-performing system in challenge)
22. C.-M. Chien, M. Zhang, J.-C. Chou, K. Livescu,
 “Few-shot spoken language understanding via joint speech-text models,”
ASRU 2023.
(Best Student Paper Award)
23. C.-I. J. Lai, F. Shi, P. Peng, Y. Kim, K. Gimpel, S. Chang, Y.-S. Chuang, S. Bhati, D. Cox, D. Harwath, Y. Zhang, K. Livescu, J. Glass
 “Audio-visual neural syntax acquisition,”
ASRU 2023.
24. S. Shon, S. Arora, C.-J. Lin, A. Pasad, F. Wu, R. Sharman, W.-L. Wu, H.-Y. Lee, K. Livescu, S. Watanabe,
 “SLUE Phase-2: A Benchmark Suite of Diverse Spoken Language Understanding Tasks,”
ACL 2023.
25. A. Pasad, B. Shi, K. Livescu,
 “Comparative layer-wise analysis of self-supervised speech models,”
ICASSP 2023.
26. S. Shon, F. Wu, K. Kim, P. Sridhar, K. Livescu, S. Watanabe,
 “Context-aware fine-tuning of self-supervised speech models,”
ICASSP 2023.

27. B. Shi, D. Brentari, G. Shakhnarovich, K. Livescu,
 “TTIC’s WMT-SLT 22 Sign Language Translation System,”
Seventh Conference on Machine Translation (WMT) 2022.
28. B. Shi, D. Brentari, G. Shakhnarovich, and K. Livescu,
 “Open-domain sign language translation learned from online video,”
EMNLP 2022.
29. S. Toshniwal, S. Wiseman, K. Livescu, and K. Gimpel,
 “Baked-in state probing,”
EMNLP 2022.
30. A. Pasad, F. Wu, S. Shon, K. Livescu, and K. J. Han,
 “On the use of external data for spoken named entity recognition,”
NAACL 2022.
31. B. Shi, D. Brentari, G. Shakhnarovich, and K. Livescu,
 “Searching for fingerspelled content in American Sign Language,”
ACL 2022.
32. H. Shi, K. Gimpel, and K. Livescu,
 “Substructure distribution projection for zero-shot cross-lingual dependency parsing,”
ACL 2022.
33. S. Shon, A. Pasad, F. Wu, P. Brusco, Y. Artzi, K. Livescu, and K. J. Han,
 “SLUE: New benchmark tasks for spoken language understanding evaluation on natural
 speech,”
ICASSP 2022.
34. S. Toshniwal, S. Wiseman, K. Livescu, and K. Gimpel,
 “Chess as a testbed for language model state tracking,”
AAAI 2022.
35. A. Pasad, J.-C. Chou, and K. Livescu,
 “Layer-wise analysis of a self-supervised speech representation model,”
ASRU 2021.
36. S. Toshniwal*, P. Xia*, S. Wiseman, K. Gimpel, and K. Livescu,
 “On generalization in coreference resolution,”
*Fourth Workshop on Computational Models of Reference, Anaphora, and Coreference
 (CRAC) 2021.*
(Best Short Paper Award)
37. B. Shi, D. Brentari, G. Shakhnarovich, and K. Livescu,
 “Fingerspelling detection in American Sign Language,”
CVPR 2021.
38. H. Shi, K. Livescu, and K. Gimpel,
 “Substructure substitution: Structured data augmentation for NLP,”
Findings of ACL-IJCNLP 2021.
39. Y. Hu, S. Settle, and K. Livescu,
 “Acoustic span embeddings for multilingual query-by-example search,”
SLT 2021.
40. B. Shi, S. Settle, and K. Livescu,
 “Whole-word segmental speech recognition with acoustic word embeddings,”
SLT 2021.

41. P. Peng, H. Kamper, and K. Livescu,
“A correspondence variational autoencoder for unsupervised acoustic word embeddings,”
NeurIPS Workshop on Self-Supervised Learning for Speech and Audio Processing 2020.
42. H. Shi, K. Livescu, and K. Gimpel,
“On the role of supervision in unsupervised constituency parsing,”
EMNLP 2020.
43. S. Toshniwal, S. Wiseman, A. Ettinger, K. Livescu, and K. Gimpel,
“Learning to ignore: Long document coreference with bounded memory neural networks,”
EMNLP 2020.
44. Y. Hu, S. Settle, and K. Livescu,
“Multilingual jointly trained acoustic and written word embeddings,”
Interspeech 2020.
45. S. Jin, S. Wiseman, K. Stratos, and K. Livescu,
“Discrete latent variable representations for low-resource text classification,”
ACL 2020.
46. S. Toshniwal, A. Ettinger, K. Gimpel, and K. Livescu,
“PeTra: A sparsely supervised memory model for people tracking,”
ACL 2020.
47. S. Toshniwal, H. Shi, B. Shi, L. Gao, K. Livescu, and K. Gimpel,
“A cross-task analysis of text span representations,”
ACL Workshop on Representation Learning for NLP (RepL4NLP) 2020.
48. W. Wang, Q. Tang, and K. Livescu,
“Unsupervised pre-training of bidirectional speech encoders via masked reconstruction,”
ICASSP 2020.
49. B. Shi, A. Martinez Del Rio, J. Keane, D. Brentari, G. Shakhnarovich, and K. Livescu,
“Fingerspelling recognition in the wild with iterative visual attention,”
ICCV 2019.
50. T. Hayashi, S. Watanabe, T. Toda, K. Takeda, S. Toshniwal, and K. Livescu,
“Pre-trained text embeddings for enhanced text-to-speech synthesis,”
Interspeech 2019.
51. A. Pasad, B. Shi, H. Kamper, and K. Livescu,
“On the contributions of visual and textual supervision in low-resource semantic speech retrieval,”
Interspeech 2019.
52. H. Shi, J. Mao, K. Gimpel, and K. Livescu,
“Visually grounded neural syntax acquisition,”
ACL 2019
(**Best Paper Nominee**).
53. S. Bansal, H. Kamper, K. Livescu, A. Lopez, and S. Goldwater,
“Pre-training on high-resource speech recognition improves low-resource speech-to-text translation,”
NAACL HLT 2019.
54. H. Kamper, A. Anastassiou, and K. Livescu,
“Semantic query-by-example speech search using visual grounding,”
ICASSP 2019.

55. S. Settle, K. Audhkhasi, K. Livescu, and M. Picheny,
“Acoustically grounded word embeddings for improved acoustics-to-word speech recognition,”
ICASSP 2019.
56. B. Shi, A. Martinez Del Rio, J. Keane, J. Michaux, D. Brentari, G. Shakhnarovich, and K. Livescu,
“American Sign Language fingerspelling recognition in the wild,”
SLT 2018.
57. S. Toshniwal, A. Kannan, C.-C. Chiu, Y. Wu, T. N. Sainath, and K. Livescu,
“A comparison of techniques for language model integration in encoder-decoder speech recognition,”
SLT 2018.
58. M. Chen, Q. Tang, K. Livescu, and K. Gimpel,
“Variational sequential labelers for semi-supervised learning,”
EMNLP 2018.
59. S. Bansal, H. Kamper, K. Livescu, A. Lopez, S. Goldwater,
“Low-resource speech-to-text translation,”
Interspeech 2018.
60. T. Tran, S. Toshniwal, M. Bansal, K. Gimpel, K. Livescu, and M. Ostendorf, “Parsing speech: A neural approach to integrating lexical and acoustic-prosodic information,”
NAACL HLT 2018.
61. K. Krishna, L. Lu, K. Gimpel, and K. Livescu,
“A Study of All-Convolutional Encoders for Connectionist Temporal Classification,”
ICASSP 2018.
62. Q. Tang, W. Wang, and K. Livescu,
“Acoustic feature learning using cross-domain articulatory measurements,”
ICASSP 2018.
63. H. Kamper, K. Livescu, and S. Goldwater,
“An embedded segmental K-means model for unsupervised segmentation and clustering of speech,”
ASRU 2017
(Best Paper Nominee).
64. B. Shi and K. Livescu,
“Multitask training with unlabeled data for end-to-end sign language fingerspelling recognition,”
ASRU 2017.
65. S. Toshniwal, H. Tang, L. Lu, and K. Livescu,
“Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recognition,”
Interspeech 2017.
66. H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu,
“Visually grounded learning of keyword prediction from untranscribed speech,”
Interspeech 2017.
67. S. Settle, K. Levin, H. Kamper, and K. Livescu,
“Query-by-example search with discriminative neural acoustic word embeddings,”
Interspeech 2017.

68. Q. Tang, W. Wang, and K. Livescu,
“Acoustic feature learning with deep variational canonical correlation analysis,”
Interspeech 2017.
69. L. Tu, K. Gimpel, and K. Livescu,
“Learning to Embed Words in Context for Syntactic Tasks,”
ACL Workshop on Representation Learning for NLP (RepL4NLP) 2017
(**Best Paper**).
70. W. He, W. Wang, and K. Livescu,
“Multi-view recurrent neural acoustic word embeddings,”
ICLR 2017.
71. S. Settle and K. Livescu,
“Discriminative acoustic word embeddings: Recurrent neural network-based approaches,”
SLT 2016.
72. S. Toshniwal and K. Livescu,
“Jointly learning to align and convert graphemes to phonemes with neural attention models,”
SLT 2016.
73. H. Tang, W. Wang, K. Gimpel, and K. Livescu,
“End-to-end training approaches for discriminative segmental models,”
SLT 2016.
74. J. Wieting, M. Bansal, K. Gimpel, and K. Livescu,
“CHARAGRAM: Embedding words and sentences via character n-grams,”
EMNLP 2016.
75. H. Tang, W. Wang, K. Gimpel, and K. Livescu,
“Efficient segmental cascades for speech recognition,”
Interspeech 2016.
76. W. Wang, H. Tang, K. Livescu,
“Triphone state-tying via deep canonical correlation analysis,”
Interspeech 2016.
77. P. Swaroop Madhyastha, M. Bansal, K. Gimpel, and K. Livescu,
“Mapping unseen words to task-trained embedding spaces,”
ACL Workshop on Representation Learning for NLP (RepL4NLP) 2016
(**Best Paper**).
78. T. Michaeli, W. Wang, and K. Livescu,
“Nonparametric canonical correlation analysis,”
ICML 2016.
79. W. Wang and K. Livescu,
“Large-scale approximate kernel canonical correlation analysis,”
ICLR 2016.
80. J. Wieting, M. Bansal, K. Gimpel, and K. Livescu,
“Towards universal paraphrastic sentence embeddings,”
ICLR 2016.
81. H. Kamper, W. Wang, and K. Livescu,
“Deep convolutional acoustic word embeddings using word-pair side information,”
ICASSP 2016.

82. T. Kim, W. Wang, H. Tang, and K. Livescu,
“Signer-independent fingerspelling recognition with deep neural network adaptation,”
ICASSP 2016
(**Best Student Paper nominee**).
83. W. Wang, R. Arora, N. Srebro, and K. Livescu,
“Stochastic optimization for deep CCA via nonlinear orthogonal iterations,”
Allerton Conference on Communication, Control, and Computing 2015.
84. H. Tang, W. Wang, K. Gimpel, and K. Livescu,
“Discriminative segmental cascades for feature-rich phone recognition,”
ASRU 2015
(**Best Paper Nominee**).
85. W. Wang, R. Arora, K. Livescu, and J. Bilmes,
“On deep multi-view representation learning,”
ICML 2015.
86. A. Lu, W. Wang, M. Bansal, K. Gimpel, and K. Livescu,
“Deep multilingual correlation for improved word embeddings,”
NAACL 2015.
87. W. Wang, R. Arora, K. Livescu, and J. Bilmes,
“Unsupervised learning of acoustic features via deep canonical correlation analysis”,
ICASSP 2015.
88. W. Wang, R. Arora, and K. Livescu,
“Reconstruction of articulatory measurements with smoothed low-rank matrix completion,”
SLT 2014.
89. H. Tang and K. Gimpel and K. Livescu,
“A comparison of training approaches for discriminative segmental models,”
Interspeech 2014.
90. P. Jyothi and K. Livescu,
“Revisiting word neighborhoods for speech recognition,”
ACL MORPHFSM Workshop 2014.
91. M. Bansal and K. Gimpel and K. Livescu,
“Tailoring continuous word representations for dependency parsing,”
ACL 2014.
92. R. Arora and K. Livescu,
“Multi-view learning with supervision for transformed bottleneck features,”
ICASSP 2014.
93. K. Levin, K. Henry, A. Jansen, and K. Livescu,
“Fixed-dimensional acoustic embeddings of variable-length segments in low-resource settings,”
ASRU 2013
(**Best Student Paper 2nd Place**)
94. T. Kim, G. Shakhnarovich, and K. Livescu,
“Fingerspelling recognition with semi-Markov conditional random fields,”
ICCV 2013.

95. P. Jyothi, E. Fosler-Lussier, and K. Livescu,
 “Discriminative training of WFST factors with application to pronunciation modeling,”
Interspeech 2013.
96. G. Andrew, R. Arora, J. Bilmes, and K. Livescu,
 “Deep canonical correlation analysis,”
ICML 2013.
(Test of Time Award Runner-Up)
97. R. Prabhavalkar, K. Livescu, E. Fosler-Lussier, and J. Keshet,
 “Discriminative articulatory models for spoken term detection in low-resource conversational settings,”
ICASSP 2013.
98. R. Arora and K. Livescu,
 “Multi-view CCA-based acoustic features for phonetic recognition across speakers and domains,”
ICASSP 2013.
99. T. Kim, K. Livescu, and G. Shakhnarovich,
 “American Sign Language fingerspelling recognition with phonological feature-based tandem models,”
SLT 2012.
100. R. Arora, A. Cotter, K. Livescu, and N. Srebro,
 “Stochastic optimization for PCA and PLS,”
Allerton Conference on Communication, Control, and Computing 2012.
101. R. Arora and K. Livescu,
 “Kernel CCA for multi-view learning of acoustic features using articulatory measurements,”
Symposium on Machine Learning in Speech and Language Processing (MLSLP) 2012.
102. R. Prabhavalkar, J. Keshet, K. Livescu, and E. Fosler-Lussier,
 “Discriminative spoken term detection with limited data,”
Symposium on Machine Learning in Speech and Language Processing (MLSLP) 2012.
103. P. Jyothi, E. Fosler-Lussier, and K. Livescu,
 “Discriminatively learning factorized finite state pronunciation models from dynamic Bayesian networks,”
Interspeech 2012.
(Best Student Paper)
104. H. Tang, J. Keshet, and K. Livescu,
 “Discriminative pronunciation modeling: A large-margin, feature-rich approach,”
ACL 2012.
105. S. Bharadwaj, R. Arora, K. Livescu, and M. Hasegawa-Johnson,
 “Multi-view acoustic feature learning using articulatory measurements,”
IEEE International Workshop on Statistical Machine Learning for Speech Processing (IWSML) 2012.
106. R. Prabhavalkar, E. Fosler-Lussier, and K. Livescu,
 “A factored conditional random field model for articulatory feature forced transcription,”
ASRU 2011.

107. J. Labiak and K. Livescu,
“Nearest neighbors with learned distances for phonetic frame classification,”
Interspeech 2011.
108. A. B. Næss, K. Livescu, and R. Prabhavalkar,
“Articulatory feature classification using nearest neighbors,”
Interspeech 2011.
109. P. Jyothi, K. Livescu, and E. Fosler-Lussier,
“Lexical access experiments with context-dependent articulatory feature-based models,”
ICASSP 2011.
110. S. Bowman and K. Livescu,
“Modeling pronunciation variation with context-dependent articulatory feature decision trees,”
Interspeech 2010.
111. L. Terry, K. Livescu, J. Pierrehumbert, and A. Katsaggelos,
“Audio-visual anticipatory coarticulation modeling by human and machine,”
Interspeech 2010.
112. A. Margolis, K. Livescu, and M. Ostendorf,
“Semi-supervised domain adaptation for automatic dialog act tagging,”
ACL Workshop on Domain Adaptation for Natural Language Processing (DANLP) 2010.
113. A. Margolis, M. Ostendorf, and K. Livescu,
“Cross-genre training for automatic prosody classification,”
Speech Prosody 2010.
114. K. Livescu and M. Stoehr,
“Multi-view learning of acoustic features for speaker recognition,”
ASRU 2009.
115. K. Chaudhuri, S. Kakade, K. Livescu, and K. Sridharan,
“Multi-view clustering via canonical correlation analysis,”
ICML 2009.
116. K. Livescu, B. Zhu, and J. Glass,
“On the phonetic information in ultrasonic microphone signals,”
ICASSP 2009.
117. O. Cetin, M. Magimai-Doss, K. Livescu, A. Kantor, S. King, C. Bartels, and J. Frankel,
“Monolingual and crosslingual comparison of tandem features derived from articulatory and phone MLPs,”
ASRU 2009.
118. J. Frankel, M. Magimai-Doss, S. King, K. Livescu, and O. Cetin,
“Articulatory feature classifiers trained on 2000 hours of telephone speech,”
Interspeech 2007.
119. M. Hasegawa-Johnson, K. Livescu, P. Lal, and K. Saenko,
“Audiovisual speech recognition with articulator positions as hidden variables,”
International Congress on Phonetic Sciences (ICPhS) 2007.
120. K. Livescu, O. Cetin, M. Hasegawa-Johnson, S. King, C. Bartels, N. Borges, A. Kantor, P. Lal, L. Yung, A. Bezman, S. Dawson-Haggerty, B. Woods, J. Frankel, M. Magimai-Doss, and K. Saenko,
“Articulatory feature-based methods for acoustic and audio-visual speech recognition:

Summary from the 2006 JHU Summer Workshop,”
ICASSP 2007.

121. O. Cetin, A. Kantor, S. King, C. Bartels, M. Magimai-Doss, J. Frankel, and K. Livescu,
“An articulatory feature-based tandem approach and factored observation modeling,”
ICASSP 2007.
122. K. Livescu, A. Bezman, N. Borges, L. Yung, O. Cetin, J. Frankel, S. King, M. Magimai-Doss, X. Chi, and L. Lavoie,
“Manual transcription of conversational speech at the articulatory feature level,”
ICASSP 2007.
123. K. Saenko and K. Livescu,
“An Asynchronous DBN for Audio-Visual Speech Recognition,”
SLT 2006.
124. K. Saenko, K. Livescu, M. Siracusa, K. Wilson, J. Glass, and T. Darrell,
“Visual speech recognition with loosely synchronized feature streams,”
ICCV 2005.
125. M. Hasegawa-Johnson, J. Baker, S. Borys, K. Chen, E. Coogan, S. Greenberg, A. Juneja, K. Kirchhoff, K. Livescu, K. Sonmez, S. Mohan, J. Muller, and T. Wang,
“Landmark-based speech recognition: Report of the 2004 Johns Hopkins Summer Workshop,”
ICASSP 2005.
126. K. Saenko, K. Livescu, J. Glass, and T. Darrell,
“Production domain modeling of pronunciation for visual speech recognition,”
ICASSP 2005.
127. K. Livescu and J. Glass,
“Feature-based pronunciation modeling with trainable asynchrony probabilities,”
ICSLP 2004.
128. K. Livescu and J. Glass,
“Feature-based pronunciation modeling for speech recognition,”
HLT/NAACL 2004.
129. K. Livescu, J. Glass, and J. Bilmes,
“Hidden feature models for speech recognition using dynamic Bayesian networks,”
Eurospeech 2003.
130. T. J. Hazen, I. Lee Hetherington, H. Shu, and K. Livescu,
“Pronunciation modeling using a finite-state transducer representation,”
ISCA Workshop on Pronunciation Modeling and Lexicon Adaptation (PMLA) 2002.
131. G. Zweig, J. Bilmes, T. Richardson, K. Filali, K. Livescu, P. Xu, K. Jackson, Y. Brandman, E. Sandness, E. Holtz, J. Torres, B. Byrne,
“Structurally discriminative graphical models for automatic speech recognition: Results from the 2001 Johns Hopkins Summer Workshop,”
ICASSP 2002.
132. K. Livescu and J. Glass,
“Segment-based recognition on the PhoneBook task: Initial results and observations on duration modeling,”
Eurospeech 2001.

133. K. Livescu and J. Glass,
“Lexical modeling of non-native speech for automatic speech recognition,”
ICASSP 2000.

TEACHING

- ◇ **TTI-Chicago/U. Chicago**, 2017–2023
TTIC 31220: Unsupervised Learning and Data Analysis
- ◇ **TTI-Chicago/U. Chicago**, 2009–2025
TTIC 31110: Speech Technologies
- ◇ **TTI-Chicago/U. Chicago**, 2011–2015
TTIC 31090: Signals, Systems, and Random Processes
- ◇ **Weizmann Institute**, Winter 2011
20114231: Introduction to Speech Recognition
- ◇ **EECS Department, MIT**, Spring 2007
Co-lecturer, 6.345: Automatic Speech Recognition
- ◇ **EECS Department, MIT**, 2005–2007
Recitation instructor, 6.003: Signals and Systems.

ADVISING (CURRENT)

PH.D.

Marcelo Beramendi Caballero, 2024–present
Tianyang (Sally) Xu, 2024–present
Jessica Yang, 2024–present
Chung-Ming Chien, 2022–present
Ju-Chieh Chou, 2021–present
Ankita Pasad, 2017–present

ADVISING (PAST)

POST-DOC

Herman Kamper 2017 (→ Assistant Professor, Stellenbosch U.)
Weiran Wang 2014–16 (→ Google)
Raman Arora 2011–13 (→ Assistant Professor, JHU)

PH.D.

Freda Shi, TTIC Ph.D. 2024 (→ Assistant Professor at U. Waterloo)
Shane Settle, TTIC Ph.D. 2023 (→ Google)
Bowen Shi, TTIC Ph.D. 2023 (→ Meta)
Qingming Tang, TTIC Ph.D. 2023 (→ Amazon)
Shubham Toshniwal, TTIC Ph.D. 2022 (→ Meta → NVIDIA)
Hao Tang, TTIC Ph.D 2017 (→ post-doc at MIT; now Lecturer at U. Edinburgh)

Taehwan Kim, TTIC Ph.D. 2016 (→ post-doc at Caltech; now Assistant Professor at UNIST)

Bahador Nooraei, TTIC 2012–15 (left program → Groupon)

Arild Næss, NTNU Ph.D 2014 (co-advised, primary advisor: Torbjørn Svendsen; → Associate Professor at Trondheim Business School)

Louis Terry, Northwestern U. Ph.D. 2011 (co-advised, primary advisor: Aggelos Katsaggelos)

M.S.

Puyuan (Jason) Peng, U. Chicago MS 2020 (co-advised, primary advisor: Mei Wang; → Ph.D. program at U. T. Austin)

John Labiak, U. Chicago MS 2010 (co-advised, primary advisor: Yali Amit; → Two Sigma)

Bo Zhu, MIT MSE 2008 (co-advised, primary advisor: James Glass; → Ph.D. program at MIT)

VISITING STUDENTS

Shester Gueuwou, 2023–present (Kwame Nkrumah University of Science and Technology, Ghana)

Yanhong Li, 2023–present (U. Chicago)

Songcheng Cai, 2023–2024 (Zhejiang U.)

Puyuan (Jason) Peng, 2021 (U. Chicago → Ph.D. student at U. T. Austin)

Yushi Hu, 2019–2021 (U. Chicago → Ph.D. student at U. Washington)

Shuning Jin, 2019–2020 (U. Minnesota Duluth → Ph.D. student at Rutgers U.)

Jack Huang, 2018 (undergrad at U. Chicago)

Yang Chen, 2018 (M.S. student at U. Chicago)

Sameer Bansal, 2018 (Ph.D. student at U. Edinburgh)

Jon Michaux, 2017 (Ph.D. at U. Chicago)

Kalpesh Krishna, 2017 (B.S. student at IIT Bombay → Ph.D. student at U. Mass. Amherst)

Wanjia He, 2016 (M.S. student at U. Chicago → Google)

Herman Kamper, 2015 and 2016 (Ph.D. student at U. Edinburgh → faculty at Stellenbosch)

Raci Lynch, 2016 (undergrad at Stanford U.)

Deblin Bagchi, 2016 (Ph.D. student at Ohio State U.)

Trang Tran, 2016 (Ph.D. student at U. Washington)

Dhivya Eswaran, 2015 (Ph.D. student at CMU)

John Wieting, 2014 (Ph.D. student at UIUC → Ph.D. student at CMU)

Pranava Swaroop Madhyastha, 2014 (Ph.D. student at UPC Barcelona)

Ang Lu, 2014 (undergrad at Tsinghua U. → M.S. student at CMU)

Hadas Benisty, 2013 (Ph.D. student at Technion → post-doc at Technion)

Victoria Evelkin, 2012 (undergrad at Technion → Ph.D. student at U. Washington)
 Sujeeth Bharadwaj, 2011 (Ph.D. student at UIUC → Google)
 Preethi Jyothi, 2010 (Ph.D. student at Ohio State U. → Assistant Professor at IIT Bombay)
 Gabrielle Knight, 2010 (undergrad at Northwestern U. → Google)
 Rohit Prabhavalkar, 2010 (Ph.D. student at Ohio State U. → Google)
 Anna Margolis, 2009–10 (Ph.D. student at U. Washington → Nuance)
 Sam Bowman, 2009 (undergrad at U. Chicago → Assistant Professor at NYU)
 Mark Stoehr, 2009 (undergrad at U. Chicago → Ph.D. student at U. Chicago)

THESIS COMMITTEES

Yifan Peng, CMU, expected 2025
 Pero Savarese, TTIC, expected 2025
 Lingyu Gao, TTIC, 2024
 Nay San, Stanford, 2024
 Salah Zaiem, Institut Polytechnique de Paris, 2024
 Robin Algayres, Ecole Normale Supérieure/INRIA, 2023
 Davis Yoshida, TTIC, 2023
 Lifu Tu, TTIC, 2021 (→ post-doc, Salesforce Research)
 Mohammadreza Mostajabi, TTIC, 2019
 Bowen Wang, U. Chicago M.S., 2018
 Jianzhu Ma, TTIC, 2015 (→ post-doc, UCSD)
 Jonathan Keane, U. Chicago, 2015 (→ post-doc, U. Chicago)
 Line Adde, NTNU, 2013 (→ Max Manus)
 Preethi Jyothi, Ohio State U., 2013 (→ Assistant Professor at IIT Bombay)
 Arthur Kantor, UIUC, 2013 (→ IBM)
 Rohit Prabhavalkar, Ohio State U., 2013 (→ Google)
 Sravana Reddy, U. Chicago, 2012 (→ Hess Fellow at Wellesley)
 Sonija Waxmonsky, U. Chicago, 2011 (→ LexisNexis)
 James Kirby, U. Chicago, 2010 (→ Lecturer at U. Edinburgh)
 Siwei Wang, U. Chicago, 2010 (→ post-doc at Argonne National Lab)

SERVICE

PROGRAM COMMITTEES

Technical/program co-chair: Interspeech 2022; ICLR 2019; ASRU 2015, 2017, 2019, 2025
Area chair: ARR 2025; Interspeech 2020, 2021; EMNLP 2018, 2020; ICLR 2018, 2020; ICASSP 2015–18, 2025

CONFERENCE/WORKSHOP ORGANIZATION

Organizer/organizing committee member: TTIC Summer Workshop on the Foundations of Speech and Audio Foundation Models 2025; ASRU SPARKS Workshop on Speech Foundation Models and their Performance Benchmarks 2023; AAAI Workshop on Self-Supervised Learning for Audio and Speech Processing 2022; NeurIPS Workshop on Self-Supervised Learning for Speech and Audio Processing 2020; Workshop for Young Female Researchers in Speech Science & Technology 2016–19; ISCA Workshop on Machine Learning in Speech and Language Processing 2012, 2016–2017; ICML Workshop on Multi-View Representation Learning 2016; Midwest Speech and Language Days 2013, 2015, 2017, 2019; Workshop on Speech Production in Automatic Speech Recognition, 2013; Speech Prosody 2010; Illinois Speech Day 2009–10

Workshops chair: EMNLP 2013

EDITING AND REVIEWING

Associate/subject editor: *Transactions of the ACL*, 2021–present; *IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI)*, 2020–2024; *IEEE Open Journal of Signal Processing (OJSP)* 2020–2021; *IEEE Trans. Audio, Speech, and Language Processing (TASLP)*, 2014–17; *Speech Communication*, 2012–13

Guest editor: *IEEE JSTSP* Special Issue on Self-Supervised Learning for Speech and Audio Processing, 2022; Special issue of *Computer Speech and Language* on Speech Production in Speech Technologies, 2016

Reviewer: *IEEE Trans. Sig. Proc.*, *IEEE Trans. Multimedia*, *IEEE Trans. Audio, Speech and Lang. Proc.*, *Journal of Machine Learning Research (JMLR)*, *Speech Communication*, *Computer Speech & Language*, Interspeech, ICASSP, ASRU, SLT, NeurIPS, ICML, ICLR, AISTATS, CVPR, ACL, EACL, HLT/NAACL, SIGMORPHON

Panelist: NSF Robust Intelligence programs, 2009, 2013, 2016, 2019, 2021, 2024

External reviewer: NSF IIS, 2022–2024; NSF BCS, 2017

OTHER PROFESSIONAL SERVICE

Panelist: U. Chicago Rising Stars Workshop, 2021, 2022; ISCA Doctoral Consortium, 2017; ISCA Student-meet-Expert Event, 2016

Member: IEEE Speech and Language Technical Committee, 2011–17, 2024–present

Secretary/Workshop Coordinator: ISCA Special Interest Group on Machine Learning (SIGML), 2011–17

Tutorial co-presenter: NAACL 2022 & Interspeech 2022 (topic: Self-Supervised Representation Learning for Speech Processing); SLT 2014 (topic: Multi-View Learning of Representations for Speech and Language)

INSTITUTIONAL SERVICE

TTIC Faculty Hiring Committee chair, 2022–23, 2024–2025

TTIC Faculty Hiring Committee co-chair, 2021–22

TTIC Colloquium Coordinator, 2008–09, 2014–23

TTIC Student Support Coordinator, 2018–23

Faculty advisor, TTIC Student Workshop, 2021–23

Research Advisory Committee, U. Chicago Data Science Institute, 2021–23

Committee member, TTIC Renovation Design Committee, 2019–20

Organizing committee member, TTIC Student Workshop, 2016–20

Steering committee member, U. Chicago Center for Data and Computing, 2019–2021

Committee member, TTIC Presidential Search Support Team, 2018

Committee member, TTIC Committee on Discrimination, Harassment and Abusive Behavior policy, 2017

Committee member, TTIC Website Committee, 2016–20

Committee member, TTIC Computing Committee, 2010–11

Research faculty mentor to Kevin Gimpel (2013–15; now Assistant Professor at TTIC), Mohit Bansal (2013–16; now Professor at UNC Chapel Hill), Liang Lu (2016–17; now at MSR), Sam Wiseman (2018–21; now Assistant Professor at Duke), Kartik Goyal (2022–23; now Assistant Professor at Georgia Tech), Zhewei Sun (2024–present)

SELECTED INVITED TALKS

WASPAA (**keynote speaker**), 2025

Midwest Speech and Language Days (**keynote speaker**), 2025

Allen Institute for AI, 2024

Forum for AI, UT Austin, 2024

Linguistics Colloquium, Stanford University, 2024

Google Research, 2024

CLSP Seminar, Johns Hopkins University, 2023

LTI Colloquium, Carnegie Mellon University, 2023

ICASSP Perspectives Talk, Rhodes, Greece, 2023

ICASSP Workshop on Self-Supervision in Audio, Speech and Beyond, Rhodes, Greece, 2023

International School on Deep Learning (DeepLearn) (**invited instructor**), Bournemouth, 2023

NeurIPS Workshop on Self-Supervised Learning: Theory and Practice, 2022

FAIR Seminar, Meta, 2022

U. Cambridge LTL Seminar, 2022

Meta Speech & Audio Summit, 2021

U. Virginia ECE Seminar Series, 2021

U. T. Austin Forum for AI, 2020

Amazon AWS, 2020

ICML Workshop on Self-Supervision in Audio and Speech, 2020

Seminar @ Cornell Tech, New York, 2019

New York Academy of Science Symposium on Natural Language, Dialog, and Speech (**keynote speaker**), New York, 2019

Speech and Audio in the Northeast (SANE) Workshop, New York, 2019

Auditory EEG Signal Processing (AESoP) symposium, Leuven, Belgium, 2019

Machine Learning Summer School (**invited instructor**), London, 2019
IEEE Workshop on Spoken Language Technology, Athens, 2018
LTI Colloquium, Carnegie Mellon University, 2018
Google Multimodal Machine Perception Workshop, San Francisco, 2018
Lisbon Machine Learning School, Lisbon, 2018
ML4Audio NeurIPS Workshop, Long Beach, CA, 2017
AI Seminar, U. Michigan, 2017
EMNLP Workshop on Subword & Character Level Models in NLP, Copenhagen, 2017
Workshop on Deep Learning, U. Maryland, 2017
AAAI Spring Symposium on the Science of Intelligence (**keynote speaker**), Stanford University, 2017
Nuance Research Conference (**keynote speaker**), 2017
Workshop on Speech Representation, Perception, and Recognition, McGovern Institute, MIT, 2017
Symposium on Speech and Audio Processing, ICSEE, Israel, 2016
Jelinek Summer Workshop School on Human Language Technology (**tutorial instructor**), 2015–2016