

Karen Livescu

Associate 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.

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

◇ **Associate Professor**, Toyota Technological Institute at Chicago, 2017–present

◇ **Associate Professor, part-time**, Computer Science Department, U. Chicago, 2017–present

◇ **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

◇ **Amazon AWS Machine Learning Research Award**, 2020

◇ **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/Johns Hopkins Workshop student grant winner**, 2001–02
- ◇ **NSF Graduate Research Fellowship**, 1997–2000

EXTERNAL RESEARCH FUNDING

Research at TTIC is funded through a combination of external grants and internal institute funding. Roughly 3/4 of my Ph.D. student funding has come from external grants, with the remainder from TTIC funding. Total external funding 2009–2020: \$3.6+ million from a mix of federal and industrial sources.

- ◇ **U. Chicago Center for Data and Computation Data Science Discovery Grant:** “A Data Processing Pipeline to Transcribe Broadcast Police Communications for Further Study.” 2019. PI: Margaret Beale Spencer. Co-PI: Karen Livescu.
- ◇ **Amazon AWS Machine Learning Research Award:** “Multilingual Acoustic-Semantic Embeddings of Spoken Language.” 2019. PI: Karen Livescu.
- ◇ **Air Force Office of Sponsored Research Center of Excellence:** “Machines, Algorithms, and Data Lab (MADLab): A University Center of Excellence in Efficient and Robust Machine Learning.” 5/1/18 - 4/30/23. Center based at U. Wisconsin Madison, PI Rob Nowak. TTIC PI: Greg Shakhnarovich. Co-PI: Karen Livescu.
- ◇ **NSF IIS-1816627:** “RI: Small: From acoustics to semantics: Embedding speech for a hierarchy of tasks.” 8/15/18 - 7/31/21. PI: Karen Livescu.
- ◇ **Google Research Award:** “Discriminative Neural Segmental Models for Sequence Prediction.” 2016. PI: Karen Livescu.
- ◇ **Google Research Award:** “Spoken language understanding with multi-view learning.” 2014. PI: Mohit Bansal. Co-PIs: Kevin Gimpel, Karen Livescu.
- ◇ **NSF IIS-1409837:** RI: Medium: Collaborative Research: Models of Handshape Articulatory Phonology for Recognition and Analysis of American Sign Language. 6/1/14 - 5/31/18. PI: Karen Livescu. Co-PI: Greg Shakhnarovich (TTIC), Diane Brentari (U. Chicago), Jason Riggle (U. Chicago).
- ◇ **NSF IIS-1321015:** RI: Small: Multi-View Learning of Acoustic Features for Speech Recognition Using Articulatory Measurements. 9/1/13 - 8/31/17. PI: Karen Livescu.
- ◇ **NSF IIS-1433485:** EAGER: Discovery of Segmental Sub-Word Structure in Speech. 3/1/14 - 2/28/15. PI: Karen Livescu.

- ◇ **NSF IIS-0905633**: RI: Medium: Collaborative Research: Explicit Articulatory Models of Spoken Language, with Application to Automatic Speech Recognition. 7/1/09 - 6/30/13. PI: Karen Livescu. Co-PIs: Jeff Bilmes (U. Washington), Eric Fosler-Lussier (Ohio State U.).

PUBLICATIONS

JOURNALS

1. 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.
2. 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.
3. 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.
4. 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.
5. K. Livescu, P. Jyothi, and E. Fosler-Lussier,
“Articulatory feature-based pronunciation modeling,”
Computer Speech and Language **36**:212–232, March 2016.
6. 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.
7. K. Livescu, E. Fosler-Lussier, and F. Metze,
“Subword modeling for automatic speech recognition,”
Signal Processing Magazine **29**(6):44–57, November 2012.
8. 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.
9. 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.
10. 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. Wang, Q. Tang, and K. Livescu,
“Unsupervised pre-training of bidirectional speech encoders via masked reconstruction,”
ICASSP 2020.
2. 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.
3. 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.
4. 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.
5. H. Shi, J. Mao, K. Gimpel, and K. Livescu,
“Visually grounded neural syntax acquisition,”
ACL 2019
(**Best Paper Nominee**).
6. 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.
7. H. Kamper, A. Anastassiou, and K. Livescu,
“Semantic query-by-example speech search using visual grounding,”
ICASSP 2019.
8. S. Settle, K. Audhkhasi, K. Livescu, and M. Picheny,
“Acoustically grounded word embeddings for improved acoustics-to-word speech recogni-
tion,”
ICASSP 2019.
9. 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.
10. 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.
11. M. Chen, Q. Tang, K. Livescu, and K. Gimpel,
“Variational sequential labelers for semi-supervised learning,”
EMNLP 2018.
12. S. Bansal, H. Kamper, K. Livescu, A. Lopez, S. Goldwater,
“Low-resource speech-to-text translation,”
Interspeech 2018.
13. 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.

14. K. Krishna, L. Lu, K. Gimpel, and K. Livescu,
“A Study of All-Convolutional Encoders for Connectionist Temporal Classification,”
ICASSP 2018.
15. Q. Tang, W. Wang, and K. Livescu,
“Acoustic feature learning using cross-domain articulatory measurements,”
ICASSP 2018.
16. 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**).
17. B. Shi and K. Livescu,
“Multitask training with unlabeled data for end-to-end sign language fingerspelling recog-
nition,”
ASRU 2017.
18. S. Toshniwal, H. Tang, L. Lu, and K. Livescu,
“Multitask learning with low-level auxiliary tasks for encoder-decoder based speech recog-
nition,”
Interspeech 2017.
19. H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu,
“Visually grounded learning of keyword prediction from untranscribed speech,”
Interspeech 2017.
20. S. Settle, K. Levin, H. Kamper, and K. Livescu,
“Query-by-example search with discriminative neural acoustic word embeddings,”
Interspeech 2017.
21. Q. Tang, W. Wang, and K. Livescu,
“Acoustic feature learning with deep variational canonical correlation analysis,”
Interspeech 2017.
22. 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**).
23. W. He, W. Wang, and K. Livescu,
“Multi-view recurrent neural acoustic word embeddings,”
ICLR 2017.
24. S. Settle and K. Livescu,
“Discriminative acoustic word embeddings: Recurrent neural network-based approaches,”
SLT 2016.
25. S. Toshniwal and K. Livescu,
“Jointly learning to align and convert graphemes to phonemes with neural attention
models,”
SLT 2016.
26. H. Tang, W. Wang, K. Gimpel, and K. Livescu,
“End-to-end training approaches for discriminative segmental models,”
SLT 2016.

27. J. Wieting, M. Bansal, K. Gimpel, and K. Livescu,
“CHARAGRAM: Embedding words and sentences via character n-grams,”
EMNLP 2016.
28. H. Tang, W. Wang, K. Gimpel, and K. Livescu,
“Efficient segmental cascades for speech recognition,”
Interspeech 2016.
29. W. Wang, H. Tang, K. Livescu,
“Triphone state-tying via deep canonical correlation analysis,”
Interspeech 2016.
30. 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**).
31. T. Michaeli, W. Wang, and K. Livescu,
“Nonparametric canonical correlation analysis,”
ICML 2016.
32. W. Wang and K. Livescu,
“Large-scale approximate kernel canonical correlation analysis,”
ICLR 2016.
33. J. Wieting, M. Bansal, K. Gimpel, and K. Livescu,
“Towards universal paraphrastic sentence embeddings,”
ICLR 2016.
34. H. Kamper, W. Wang, and K. Livescu,
“Deep convolutional acoustic word embeddings using word-pair side information,”
ICASSP 2016.
35. T. Kim, W. Wang, H. Tang, and K. Livescu,
“Signer-independent fingerspelling recognition with deep neural network adaptation,”
ICASSP 2016
(**Best Student Paper nominee**).
36. 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.
37. H. Tang, W. Wang, K. Gimpel, and K. Livescu,
“Discriminative segmental cascades for feature-rich phone recognition,”
ASRU 2015
(**Best Paper Nominee**).
38. W. Wang, R. Arora, K. Livescu, and J. Bilmes,
“On deep multi-view representation learning,”
ICML 2015.
39. A. Lu, W. Wang, M. Bansal, K. Gimpel, and K. Livescu,
“Deep multilingual correlation for improved word embeddings,”
NAACL 2015.
40. W. Wang, R. Arora, K. Livescu, and J. Bilmes,
“Unsupervised learning of acoustic features via deep canonical correlation analysis”,
ICASSP 2015.

41. W. Wang, R. Arora, and K. Livescu,
“Reconstruction of articulatory measurements with smoothed low-rank matrix completion,”
SLT 2014.
42. H. Tang and K. Gimpel and K. Livescu,
“A comparison of training approaches for discriminative segmental models,”
Interspeech 2014.
43. P. Jyothi and K. Livescu,
“Revisiting word neighborhoods for speech recognition,”
ACL MORPHFSM Workshop 2014.
44. M. Bansal and K. Gimpel and K. Livescu,
“Tailoring continuous word representations for dependency parsing,”
ACL 2014.
45. R. Arora and K. Livescu,
“Multi-view learning with supervision for transformed bottleneck features,”
ICASSP 2014.
46. 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)
47. T. Kim, G. Shakhnarovich, and K. Livescu,
“Fingerspelling recognition with semi-Markov conditional random fields,”
ICCV 2013.
48. P. Jyothi, E. Fosler-Lussier, and K. Livescu,
“Discriminative training of WFST factors with application to pronunciation modeling,”
Interspeech 2013.
49. G. Andrew, R. Arora, J. Bilmes, and K. Livescu,
“Deep canonical correlation analysis,”
ICML 2013.
50. R. Prabhavalkar, K. Livescu, E. Fosler-Lussier, and J. Keshet,
“Discriminative articulatory models for spoken term detection in low-resource conversational settings,”
ICASSP 2013.
51. R. Arora and K. Livescu,
“Multi-view CCA-based acoustic features for phonetic recognition across speakers and domains,”
ICASSP 2013.
52. T. Kim, K. Livescu, and G. Shakhnarovich,
“American Sign Language fingerspelling recognition with phonological feature-based tandem models,”
SLT 2012.
53. R. Arora, A. Cotter, K. Livescu, and N. Srebro,
“Stochastic optimization for PCA and PLS,”
Allerton Conference on Communication, Control, and Computing 2012.

54. 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.
55. 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.
56. P. Jyothi, E. Fosler-Lussier, and K. Livescu,
“Discriminatively learning factorized finite state pronunciation models from dynamic Bayesian networks,”
Interspeech 2012.
(Best Student Paper)
57. H. Tang, J. Keshet, and K. Livescu,
“Discriminative pronunciation modeling: A large-margin, feature-rich approach,”
ACL 2012.
58. 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.
59. R. Prabhavalkar, E. Fosler-Lussier, and K. Livescu,
“A factored conditional random field model for articulatory feature forced transcription,”
ASRU 2011.
60. J. Labiak and K. Livescu,
“Nearest neighbors with learned distances for phonetic frame classification,”
Interspeech 2011.
61. A. B. Næss, K. Livescu, and R. Prabhavalkar,
“Articulatory feature classification using nearest neighbors,”
Interspeech 2011.
62. P. Jyothi, K. Livescu, and E. Fosler-Lussier,
“Lexical access experiments with context-dependent articulatory feature-based models,”
ICASSP 2011.
63. S. Bowman and K. Livescu,
“Modeling pronunciation variation with context-dependent articulatory feature decision trees,”
Interspeech 2010.
64. L. Terry, K. Livescu, J. Pierrehumbert, and A. Katsaggelos,
“Audio-visual anticipatory coarticulation modeling by human and machine,”
Interspeech 2010.
65. 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.
66. A. Margolis, M. Ostendorf, and K. Livescu,
“Cross-genre training for automatic prosody classification,”
Speech Prosody 2010.

67. K. Livescu and M. Stoehr,
“Multi-view learning of acoustic features for speaker recognition,”
ASRU 2009.
68. K. Chaudhuri, S. Kakade, K. Livescu, and K. Sridharan,
“Multi-view clustering via canonical correlation analysis,”
ICML 2009.
69. K. Livescu, B. Zhu, and J. Glass,
“On the phonetic information in ultrasonic microphone signals,”
ICASSP 2009.
70. 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.
71. J. Frankel, M. Magimai-Doss, S. King, K. Livescu, and O. Cetin,
“Articulatory feature classifiers trained on 2000 hours of telephone speech,”
Interspeech 2007.
72. 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.
73. 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.
74. 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.
75. 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.
76. K. Saenko and K. Livescu,
“An Asynchronous DBN for Audio-Visual Speech Recognition,”
SLT 2006.
77. K. Saenko, K. Livescu, M. Siracusa, K. Wilson, J. Glass, and T. Darrell,
“Visual speech recognition with loosely synchronized feature streams,”
ICCV 2005.
78. 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.
79. K. Saenko, K. Livescu, J. Glass, and T. Darrell,
“Production domain modeling of pronunciation for visual speech recognition,”
ICASSP 2005.

80. K. Livescu and J. Glass,
“Feature-based pronunciation modeling with trainable asynchrony probabilities,”
ICSLP 2004.
81. K. Livescu and J. Glass,
“Feature-based pronunciation modeling for speech recognition,”
HLT/NAACL 2004.
82. K. Livescu, J. Glass, and J. Bilmes,
“Hidden feature models for speech recognition using dynamic Bayesian networks,”
Eurospeech 2003.
83. 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.
84. 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.
85. K. Livescu and J. Glass,
“Segment-based recognition on the PhoneBook task: Initial results and observations on duration modeling,”
Eurospeech 2001.
86. K. Livescu and J. Glass,
“Lexical modeling of non-native speech for automatic speech recognition,”
ICASSP 2000.

TEACHING

- ◇ **TTI-Chicago/U. Chicago**, 2017–2019
TTIC 31220: Unsupervised Learning and Data Analysis
- ◇ **TTI-Chicago/U. Chicago**, 2009–2020
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.

- Ankita Pasad TTIC, 2017–present
- Shane Settle TTIC, 2016–present

Bowen Shi TTIC, 2016–present
Haoyue “Freda” Shi TTIC, 2018–present (co-advised with Kevin Gimpel)
Qingming Tang TTIC, 2016–present
Shubham Toshniwal TTIC, 2015–present (co-advised with Kevin Gimpel)
David Yunis TTIC, 2019–present

RESEARCH INTERNS/VISITING STUDENTS

Shuning Jin, 2019–present (U. Minnesota Duluth)
Yushi Hu, 2019–present (U. Chicago)

ADVISING (PAST)

POST-DOC

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

PH.D.

Hao Tang TTIC, 2010–17 (→ post-doc at MIT)
Taehwan Kim TTIC, 2012–16 (→ post-doc at Caltech)
Bahador Nooraei, TTIC, 2012–15 (left program → Groupon)
Arild Næss, NTNU, 2010–14 (co-advised, primary advisor: Torbjørn Svendsen; → Associate Professor at Trondheim Business School)
Louis Terry, Northwestern U., 2009–11 (co-advised, primary advisor: Aggelos Katsagelos)

M.S.

John Labiak, U. Chicago, 2009–10 (co-advised, primary advisor: Yali Amit; → Two Sigma)
Bo Zhu, MIT, 2007–08 (co-advised, primary advisor: James Glass; → Ph.D. program at MIT)

RESEARCH INTERNS/VISITING STUDENTS

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

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)

SELECTED RECENT SERVICE

- ◇ **Program/technical co-chair:** Interspeech 2022; ICLR 2019; ASRU 2015, 2017, 2019
- ◇ **Area chair:** Interspeech 2020; EMNLP 2018; ICLR 2018, 2020; ICASSP 2015–18
- ◇ **Organizer/organizing committee member:** Workshop for Young Female Researchers in Speech Science & Technology 2016–19; ISCA Workshop on Machine Learning in Speech and Language Processing 2016–2017; ICML Workshop on Multi-View Representation Learning 2016; Midwest Speech and Language Days 2013, 2015, 2017, 2019
- ◇ **Associate editor:** *IEEE Trans. Audio, Speech, and Language Processing* 2014–17; *IEEE Open Journal of Signal Processing (OJSP)* 2020–present; *Speech Communication* 2012–2013

- ◇ **Guest editor:** *Computer Speech and Language* special issue on Speech Production in speech Technologies 2016
- ◇ **Standing review committee member:** *Transactions of the ACL* 2014–present
- ◇ **Reviewer:** *Speech Communication, IEEE Trans. Sig. Proc., IEEE Trans. Multimedia, IEEE Trans. Audio, Speech and Lang. Proc.*, Interspeech, ICASSP, ASRU, SLT, NeurIPS, AISTATS, CVPR, ACL, EACL, HLT/NAACL, SIGMORPHON, ICML, ICLR
- ◇ **Panelist:** NSF Robust Intelligence programs, 2009, 2013, 2016, 2019
- ◇ **External reviewer:** NSF BCS, 2017
- ◇ **Member:** IEEE Speech and Language Technical Committee, 2011–17
- ◇ **Panelist:** ISCA Doctoral Consortium, 2017; ISCA Student-meet-Expert Event, 2016
- ◇ **Tutorial co-presenter:** SLT 2014 (topic: Multi-View Learning of Representations for Speech and Language)

SELECTED RECENT TALKS

- Invited speaker,** ICML Workshop on Self-Supervision in Audio and Speech, Vienna, Austria, 2020
- Keynote speaker,** New York Academy of Science Symposium on Natural Language, Dialog, and Speech, New York, 2019
- Invited speaker,** Speech and Audio in the Northeast (SANE) Workshop, New York, 2019
- Invited speaker,** Auditory EEG Signal Processing (AESoP) symposium, Leuven, Belgium, 2019
- Invited instructor,** Machine Learning Summer School, London, 2019
- Invited speaker,** IEEE Workshop on Spoken Language Technology, Athens, 2018
- Invited speaker,** LTI Colloquium, Carnegie Mellon University, 2018
- Invited speaker,** Lisbon Machine Learning School, Lisbon, 2018
- Invited speaker,** ML4Audio NIPS Workshop, Long Beach, CA, 2017
- Invited speaker,** AI Seminar, U. Michigan, 2017
- Invited speaker,** EMNLP Workshop on Subword & Character Level Models in NLP, Copenhagen, 2017
- Invited speaker,** Workshop on Deep Learning, U. Maryland, 2017
- Keynote speaker,** AAAI Spring Symposium on the Science of Intelligence, Stanford University, 2017
- Keynote speaker,** Nuance Research Conference, 2017
- Invited speaker,** Workshop on Speech Representation, Perception, and Recognition, McGovern Institute, MIT, 2017
- Invited speaker,** Symposium on Speech and Audio Processing, ICSEE, Israel, 2016
- Invited speaker,** SILO Seminar, U. Wisconsin-Madison, 2016
- Tutorial instructor,** Jelinek Summer Workshop School on Human Language Technology, 2015–2016
- Invited speaker,** LTI Colloquium, Carnegie Mellon University, 2015
- Invited speaker,** CLIP Distinguished Colloquium, U. Maryland, 2014