

Curriculum Vitae

Joseph (Yossi) Keshet

Personal Information

Contact:

jkeshet@technion.ac.il
<https://keshet.net.technion.ac.il/>

Address:

Joseph Keshet
Office: Meyer 553
Andrew and Erna Viterbi Faculty of Electrical
and Computer Engineering
Technion - Israel Institute of Technology
Haifa 3200003, Israel

Research Interests

Deep learning; Speech and language processing.

Education

- 2002 - 2008 **Ph.D. in Computer Science**, The School of Computer Science and Engineering,
The Hebrew University, Jerusalem, Israel.
Advisor: Yoram Singer
Dissertation: *Large Margin Algorithms for Discriminative Continuous Speech.*
- 2000 - 2002 **M.Sc. in Electrical Engineering**, Department of Electrical Engineering – Systems,
Tel-Aviv University, Tel-Aviv, Israel.
Advisor: Dr. Dan Chazan, IBM Research Labs, Haifa.
Dissertation: *Stop consonant spotting in continuous speech.*
- 1991 - 1994 **B.Sc. in Electrical Engineering**, Cum laude
Department of Electrical Engineering – Systems,
Tel-Aviv University, Tel-Aviv, Israel.

Awards and Scholarships

- 2017 IEEE Senior Member.
- 2010 Founder and chair of the Machine learning special interest group (SIG)
of International Speech Communication Society (ISCA).
- 2007 PASCAL Network of Excellence, Ph.D. Label Award.
- 2004-2006 Leibniz Student Fellowship.
- 2002 IBM Fellowship.
- 2001 Israel Defense Forces (IDF) – National award for outstanding research and
development achievements.

Work Experience

- 2022 - present **Associate Professor**, Faculty of Electrical and Computer Engineering,
Technion - Israel Institute of Technology

- 2020 - present **Amazon Scholar**, Amazon, Tel-Aviv, Israel
- 2018 - 2022 **Associate Professor**, Department of Computer Science, Bar-Ilan University
- 2020 - 2021 **Founding Venture Partner**, Disruptive AI Venture Capital, Herzelia, Israel
- 2013 - 2018 **Senior Lecturer (Assistant Professor)**, Department of Computer Science, Bar-Ilan University
- 2009 - 2012 **Research Assistant Professor**, TTI-Chicago and Department of Computer Science, University of Chicago, Chicago, IL.
- 2007 - 2009 **Research Scientist (Postdoc)**, IDIAP Research Institute, Martigny, Switzerland.
- 2002 - 2007 **Research Assistant**, School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel.
- 1995 - 2002 **Head of speech processing research section**, The R&D Center, Elite unit (8200), Israel Defense Forces (IDF).
Managing 2-3 researchers. Topics mainly concern speech coding, universal speech encoder, information theoretic tools for speech coding and decoding.

Teaching Experience

Machine learning, 2014–2021
 Advance techniques in machine learning, 2013–2021
 Speech processing and recognition, 2017–2021
 Probability for Computer Science, 2014–2018.

Selected Publications

Edited Books

Joseph Keshet and Samy Bengio, Eds., *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, John Wiley & Sons, March 2009.

Book Chapters

Joseph Keshet, Subhransu Maji, Tamir Hazan, and Tommi Jaakkola, Perturbation Models and PAC-Bayesian Generalization Bounds, in *Perturbations, Optimization, and Statistics*, Tamir Hazan, George Papandreou, and Daniel Tarlow, Eds., The MIT Press, 2016

Joseph Keshet, Optimizing the Measure of Performance in Structured Prediction, in *Advanced Structured Prediction*, Sebastian Nowozin, Peter V. Gehler, Jeremy Jancsary, and Christoph H. Lampert, Eds., The MIT Press, 2014.

Joseph Keshet, Shai Shalev-Shwartz, Yoram Singer and Dan Chazan, A Large Margin Algorithm for Forced Alignment, in *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, Joseph Keshet and Samy Bengio, Eds., John Wiley & Sons, March 2009.

Joseph Keshet and Dan Chazan, A Kernel Wrapper for Phoneme Sequence Recognition, in *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, Joseph Keshet and Samy Bengio, Eds., John Wiley & Sons, March 2009.

Joseph Keshet, A Proposal for a Kernel-based Algorithm for Large Vocabulary Continuous Speech Recognition, in *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, Joseph Keshet

and Samy Bengio, Eds., John Wiley & Sons, March 2009.

David Grangier, **Joseph Keshet** and Samy Bengio, Discriminative Keyword Spotting , in *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, Joseph Keshet and Samy Bengio, Eds., John Wiley & Sons, March 2009.

Journal Papers

Adar Paz, Eshkol Rafaeli, Eran Bar-Kalifa, Eva Gilboa-Schechtman, Sharon Gannot, Bracha Laufer-Goldshtein, Shrikanth Narayanan, **Joseph Keshet**, Dana Atzil-Slonim, Intrapersonal and Interpersonal Vocal Affect Dynamics During Psychotherapy, *Journal of Consulting and Clinical Psychology*, Vol. 89, Issue 3, pp. 227, ASA, March 2021.

Matthew Goldrick, Yosi Shrem, Oriana Kilbourn-Ceron, Cristina Baus, and **Joseph Keshet**, Using automated acoustic analysis to explore the link between planning and articulation in second-language speech production, *Language, Cognition and Neuroscience*, pp. 1–16, 2020

Ariel Rosenfeld, Moshe Cohen, Sarit Kraus, and **Joseph Keshet**, Online Prediction of Time Series with Assumed Behavior, *Engineering Applications of Artificial Intelligence*, Volume 88, February 2020, 103358.

Yehoshua Dissen, Jacob Goldberger, and **Joseph Keshet**, Formant Estimation and Tracking: A Deep Learning Approach, *Journal of the Acoustical Society of America*, 145 (2), February 2019.

Jacob T. Cohen, Alma Cohen, Limor Benyamini, Yossi Adi, **Joseph Keshet**, Predicting glottal closure insufficiency using fundamental frequency contour analysis, *Head & Neck, Journal of the Sciences and Specialities of the Head and Neck*, Vol. 41, Issue 7, pp. 2324-2331, 2019

Esteban Buz, Adam Buchwald, and **Joseph Keshet**, Assessing Automatic VOT Annotation using Unimpaired and Impaired Speech, *International Journal of Speech-Language Pathology*, Vol. 20, No. 6, pp. 624-634, 2018.

Joseph Keshet, Automatic Speech Recognition: A Primer for Speech Pathology Researchers, *International Journal of Speech-Language Pathology*, Vol. 20 No. 6, pp. 599-609, 2018.

Matthew Goldrick, Rhonda McClain, Emily Cibelli, Yossi Adi , Erin Gustafson, Cornelia Moers, and **Joseph Keshet**, The influence of lexical selection disruptions on articulation, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, July 2018.

Tzeviya Fuchs and **Joseph Keshet**, Spoken Term Detection Automatically Adjusted for a Given Threshold, *IEEE Journal of Selected Topics in Signal Processing*, Dec 2017, Volume 11, Issue 8, pp. 1-8.

Yossi Adi, **Joseph Keshet**, Emily Cibelli, Erin Gustafson, Cynthia Clopper, and Matthew Goldrick, Automatic Measurement of Vowel Duration via Structured Prediction, *Journal of the Acoustical Society of America*, Volume 140, Issue 6, 2016.

Nisan Ozana, Stav Buchsbaum, Yael Bishitz, Yevgeny Beiderman, Zeev Schmilovitch, Ariel Schwarz, Amir Shemer, **Joseph Keshet**, and Zeev Zalevsky, An optical remote sensor for peanut kernel abortion classification, *Applied Optics*, Vol. 55, Issue 15, pp. 4005-4010, 2016.

Yossi Adi and **Joseph Keshet**, StructED : Risk Minimization in Structured Prediction, *Journal of Machine Learning Research* 17, 1-5, 2016.

Matthew Goldrick*, **Joseph Keshet***, Erin Gustafson, Jordana Heller, and Jeremy Needle, Automatic Analysis of Slips of the Tongue: Insights into the Cognitive Architecture of Speech Production, *Cognition*, 149, 31-39, 2016.

Morgan Sonderegger and **Joseph Keshet**, Automatic Discriminative Measurement of Voice Onset Time, *Journal of the Acoustical Society of America*, Vol. 132, Issue 6, pp. 3965–3979, 2012.

Francesco Orabona, **Joseph Keshet** and Barbara Caputo, Bounded Kernel-Based Online Learning, *Journal of Machine Learning Research*, 10(Nov): 2643–2666, 2009.

Joseph Keshet, David Grangier and Samy Bengio, Discriminative Keyword Spotting, *Speech Communication*, Volume 51, Issue 4, pp. 317-329, April 2009.

Joseph Keshet, Shai Shalev-Shwartz, Yoram Singer and Dan Chazan, Large Margin Algorithm for Speech-to-Phoneme and Music-to-Score Alignment, *IEEE Trans. on Audio, Speech and Language Processing*, 15(8), pp. 2373–2382, Nov. 2007.

Koby Crammer, Ofer Dekel, **Joseph Keshet**, Shai Shalev-Shwartz and Yoram Singer, Online Passive-Aggressive Algorithms, *Journal of Machine Learning Research*, 7(Mar): 551–585, 2006.

Refereed Conference papers

Bronya Roni Chernyak, Bhiksha Raj, Tamir Hazan, **Joseph Keshet**, Constant Random Perturbations Provide Adversarial Robustness with Minimal Effect on Accuracy, RobustML Workshop, Ninth International Conference on Learning Representations (ICLR), 2021.

Shahar Segal, Yossi Adi, Benny Pinkas, Carsten Baum, Chaya Ganesh, and **Joseph Keshet**, Fairness in the Eyes of the Data: Certifying Machine-Learning Models, AAI/ACM conference on AI, Ethics, and Society (AIES), 2021.

Tzeviya Sylvia Fuchs, Yael Segal, and **Joseph Keshet**, CNN-based Spoken Term Detection and Localization without Dynamic Programming, the 46rd IEEE International Conference in Acoustic, Speech and Signal Processing (ICASSP), 2021.

Felix Kreuk, **Joseph Keshet** and Yossi Adi, Self-Supervised Contrastive Learning for Unsupervised Phoneme Segmentation, The 21th Annual Conference of the International Speech Communication Association (Interspeech), Shanghai, China, 2020.

Felix Kreuk, Yossi Adi, Bhiksha Raj, Rita Singh and **Joseph Keshet**, Hide and Speak: Towards Deep Neural Networks for Speech Steganography, The 21th Annual Conference of the International Speech Communication Association (Interspeech), Shanghai, China, 2020.

Ben Goldberger, Guy Katz, Yossi Adi and **Joseph Keshet**, Minimal Modifications of Deep Neural Networks using Verification, 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR23), 2020.

Felix Kreuk, Yaniv Sheena, **Joseph Keshet**, Yossi Adi, Phoneme Boundary Detection using Learnable Segmental Features, The 45th International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Barcelona, 2020.

Yosi Shrem, Matthew Goldrick, and **Joseph Keshet**, Dr.VOT : Measuring Positive and Negative Voice Onset Time in the Wild , The 20th Annual Conference of the International Speech Communication Association (Interspeech), Graz, 2019.

Yael Segal, Tzeviya Sylvia Fuchs, **Joseph Keshet**, SpeechYOLO: Detection and Localization of Speech Objects, The 20th Annual Conference of the International Speech Communication Association (Interspeech), Graz, 2019.

Gabi Shalev, Yossi Adi, and **Joseph Keshet**, Out-of-Distribution Detection using Multiple Semantic Label Representations, Neural Information and Processing Systems (NeurIPS), 2018.

Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, **Joseph Keshet**, Turning Your Weakness Into a Strength: Watermarking Deep Neural Networks by Backdooring, Proc. of the 27th Usenix Security Symposium, Baltimore, MD, USA, 2018.

Felix Kreuk, Yossi Adi, Moustapha Cisse, **Joseph Keshet**, Fooling End-to-end Speaker Verification by Adversarial Examples, the 43rd IEEE International Conference in Acoustic, Speech and Signal Processing (ICASSP), Calgari, Canada, 2018.

Moustapha Cisse, Yossi Adi, Natalia Neverova, **Joseph Keshet**, Houdini: Fooling Deep Structured Prediction Models, Proc. of the 31st the Annual Conference on Neural Information Processing Systems (NIPS), 2017.

Yaniv Sheena, Míša Hejrná, Yossi Adi, **Joseph Keshet**, Automatic Measurement of Pre-aspiration, The 18th Annual Conference of the International Speech Communication Association (Interspeech), Stockholm, 2017.

Einat Naaman, Yossi Adi, **Joseph Keshet**, Learning Similarity Functions for Pronunciation Variations, The 18th Annual Conference of the International Speech Communication Association (Interspeech), Stockholm, 2017.

Yossi Adi, **Joseph Keshet**, Emily Cibelli, Matthew Goldrick, Sequence Segmentation Using Joint RNN and Structured Prediction Models, the 42st IEEE International Conference in Acoustic, Speech and Signal Processing (ICASSP), New Orleans, 2017.

Yehoshua Dissen and Joseph Keshet, Formant Estimation and Tracking using Deep Learning, Proc. of the 17th Annual Conference of the International Speech Communication Association (Interspeech), San Francisco, CA, 2016.

Yossi Adi, **Joseph Keshet**, Olga Dmitrieva, and Matt Goldrick, Automatic Measurement of Voice Onset Time and Prevoicing using Recurrent Neural Networks, Proc. of the 17th Annual Conference of the International Speech Communication Association (Interspeech), San Francisco, CA, 2016.

Ariel Rosenfeld, **Joseph Keshet**, Claudia V. Goldman, Sarit Kraus, Online Prediction of Exponential Decay Time Series with Human-Agent Application, Proc. of the 22st European Conference on Artificial Intelligence (ECAI), The Hague, Holland, 2016.

Rita Singh, **Joseph Keshet**, Eduard Hovy, Profiling hoax callers, the IEEE Symposium on Technologies for Homeland Security (HST), 1-6, 2016.

Rita Singh, **Joseph Keshet**, Deniz Gencaga, Bhiksha Raj, The Relationship of Voice Onset Time and Voice Offset Time to Physical Age, the 41st IEEE International Conference in Acoustic, Speech and Signal Processing (ICASSP), Shanghai, China, 2016.

Yossi Adi, **Joseph Keshet**, and Matt Goldrick, Vowel Duration Measurement using Deep Neural Networks, IEEE International Workshop on Machine Learning for Signal Processing (MLSP), Boston, USA, 2015

Adiel Ben-Shalom, **Joseph Keshet**, Roni Yeger-Granot, Automatic analysis of music: Performance of cantillation signs in Yemenite Jewish traditional cantillation, 9th Conference on Interdisciplinary Musicology (CIM14), Berlin, 2014.

Adiel Ben-Shalom, **Joseph Keshet**, Doron Modan, and Asher Laufer, Automatic Tools for Analyzing Spoken Hebrew, Afeka Conference for Speech Processing, Tel-Aviv, Israel, 2014.

Tamir Hazan, Subhransu Maji, **Joseph Keshet**, Tommi Jaakkola, On Sampling from the Gibbs distribution with Random Maximum A-Posteriori Perturbations, Proc. of the 27th the Annual Conference on Neural Information Processing Systems (NIPS), 2013.

Noam Peled, Moshe Bitan, **Joseph Keshet**, Sarit Kraus, How to predict human strategic decisions using facial expressions, International Joint Conference on Artificial Intelligence (IJCAI), 2013.

Rohit Prabhavalkar, Karen Livescu, Eric Fosler-Lussier, and **Joseph Keshet**, Discriminative Articulatory Models for Spoken Term Detection in Low-Resource Conversational Settings, IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2013.

Rohit Prabhavalkar, **Joseph Keshet**, Karen Livescu, and Eric Fosler-Lussier, Discriminative Spoken Term Detection with Limited Data, The 2nd Symposium on Machine Learning in Speech and Language Processing, Portland, Oregon, 2012.

Katharine Henry, Morgan Sonderegger and **Joseph Keshet**, Automatic Measurement of Positive and Negative Voice Onset Time, The 13th Annual Conference of the International Speech Communication Association (Interspeech), Portland, Oregon, 2012 (**full oral presentation, best paper track**).

Hao Tang, **Joseph Keshet**, and Karen Livescu, Discriminative pronunciation modeling: A large-margin, feature-rich approach, Proc. of the 50th Annual Meeting of the Association of Computational Linguistics (ACL), 2012 (**full oral presentation, best paper track**).

David McAllester and **Joseph Keshet**, Generalization Bounds and Consistency for Latent Structural Probit and Ramp Loss, Proc. of the 25th Annual Conference on Neural Information Processing Systems (NIPS), 2011 (**full oral presentation**).

Joseph Keshet, Chih-Chieh Cheng, Mark Stoehr, and David McAllester, Direct Error Rate Minimization of Hidden Markov Models, Proc. of the 12th Annual Conference of the International Speech Communication Association (Interspeech), Florence, Italy, 2011.

Andrew Cotter, Nathan Srebro, and **Joseph Keshet**, A GPU-Tailored Approach for Training Kernelized SVMs, Proc. of the 17th ACM Conference on Knowledge Discovery and Data Mining (KDD), San Diego, CA, 2011.

Joseph Keshet, David McAllester, and Tamir Hazan, PAC-Bayesian Approach for Minimization of Phoneme Error Rate, Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Prague, Czech Republic, 2011.

David McAllester, Tamir Hazan and **Joseph Keshet**, Direct Loss Minimization for Structured Prediction, Proc. of the 24th Annual Conference on Neural Information Processing Systems (NIPS), 2010.

Morgan Sonderegger and **Joseph Keshet**, Automatic Discriminative Measurement of Voice Onset Time, Proc. of the 12th Annual Conference of the International Speech Communication Association (INTERSPEECH), Chiba, Japan, 2010.

Martin Wöllmer, Florian Eyben, **Joseph Keshet**, Alex Graves, Björn Schuller, and Gerhard Rigoll, Robust Discriminative Keyword Spotting for Emotionally Colored Spontaneous Speech using Bidirectional LSTM Networks, Proc. of IEEE International Conference on Acoustic, Speech, and Signal Processing (ICASSP), Taipei, Taiwan, 2009.

Yves Grandvalet, **Joseph Keshet**, Alain Rakotomamonjy, and Stephane Canu, Support Vector Machines with a Reject Option Proc. of the 22nd Annual Conference on Neural Information Processing Systems (NIPS), 2008.

Francesco Orabona, **Joseph Keshet** and Barbara Caputo, The Projectron: a Bounded Kernel-Based Perceptron Proc. of the 25th International Conference on Machine Learning (ICML), 2008.

Joseph Keshet, David Grangier and Samy Bengio, Discriminative Keyword Spotting, Proc. of the 3rd Non-Linear Speech Processing Workshop (NOLISP), Paris, 2007.

Joseph Keshet, Shai Shalev-Shwartz, Samy Bengio, Yoram Singer and Dan Chazan, Discriminative Kernel-Based Phoneme Sequence Recognition, Proc. of the 9th International Conference on Spoken Language Processing (Interspeech), Pittsburgh, PA, 2006.

Joseph Keshet, Shai Shalev-Shwartz, Yoram Singer and Dan Chazan, Phoneme Alignment Based on Discriminative Learning, Proc. of the 9th European Conference on Speech Communication and Technology (Interspeech), Lisbon, 2005.

Shai Shalev-Shwartz, **Joseph Keshet** and Yoram Singer, Learning to Align Polyphonic Music, Proc. of the 5th International Conference on Music Information Retrieval (ISMIR), Barcelona, Spain, 2004.

Ofer Dekel, **Joseph Keshet** and Yoram Singer, An Online Algorithm for Hierarchical Phoneme Classification, Workshop on Multimodal Interaction and Related Machine Learning Algorithms (MLMI), Martigny, Switzerland, 2004, LNCS, Vol. 3361/2005, p. 146, Springer-Verlag.

Ofer Dekel, **Joseph Keshet** and Yoram Singer, Large Margin Hierarchical Classification, Proc. of the 21st International Conference on Machine Learning (ICML), Banff, Canada, 2004.

Koby Crammer, **Joseph Keshet** and Yoram Singer, Kernel Design using Boosting, in Proc. of 16th Annual Conference on Neural Information Processing Systems (NIPS), 2002.

Joseph Keshet, Dan Chazan and Ben-Zion Bobrovsky, Plosive Spotting with Large Margin Classifiers, in Proc. of 7th European Conference on Speech Communication and Technology (Eurospeech), 2001.

Professional Services

Journal Refereeing

Speech Communication, IEEE Trans. on Speech and Audio Processing, IEEE Trans. on Neural Network and Learning Systems, IEEE Journal of Selected Topics in Signal Processing, EURASIP Journal on Audio, Speech, and Music Processing, Journal of Machine Learning Research (JMLR), Pattern Recognition, Pattern Recognition Letters,

Conference Refereeing and Program Committee

Senior Area Chair in International Joint Conference on Artificial Intelligence (IJCAI), Neural Information Processing Systems (NIPS), International Conference on Machine Learning (ICML), IEEE International Conference on Acoustic, Speech and Signal Processing (ICASSP), Interspeech, North American

Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL HLT)

Other Academic Activities

Tutorial on Generating adversarial examples for speech and speaker recognition and other systems, Interspeech 2019 (with Bhiksha Raj)

Organized with Benny Pinkas a second workshop “Hacking Deep Learning²”, Deep Learning and Security, Attacking deep networks by adversarial examples, differential privacy, and forensics, Bar-Ilan University, January 22, 2019.

Organized with Benny Pinkas a workshop “Hacking Deep Learning”, Deep Learning and Security, Attacking deep networks by adversarial examples, differential privacy, and forensics, Bar-Ilan University, January 29, 2018.

Chair the third symposium on Machine Learning for Speech and Language Processing, San Francisco, CA, September 13, 2016.

<http://ttic.uchicago.edu/~klivescu/MLSLP2016>

Chair the second symposium on Machine Learning for Speech and Language Processing, Portland, OG, September 14, 2011.

<http://www.ttic.edu/sigml/symposium2012>

Chair the first symposium on Machine Learning for Speech and Language Processing, Bellevue, WA, June 27, 2011.

<http://www.ttic.edu/sigml/symposium2011>

Founder and chair of the Machine Learning for Speech and Language Processing Special Interest Group (SIGML) of International Speech Communication Association (ISCA), 2011.

Organizer of Illinois Speech Day 2012.

http://cs.biu.ac.il/~jkeshet/Illinois_Speech_Day_2012.html

Organizer of Illinois Speech Day 2011.

http://cs.biu.ac.il/~jkeshet/Illinois_Speech_Day_2011.html

Member of the IEEE Machine Learning for Signal Processing technical committee since 2011.

Organizer of the Learning Club at the Hebrew University, 2006-2007.

Student advisory committee of the International Speech Communication Association (ISCA), 2006-2007.