Curriculum Vitae
Joseph (Yossi) Keshet

Date: April 7, 2026

Personal Details

Office phone: +972 73 3787139

Email: jkeshet@technion.ac.il
WWW: https://keshet.technion.ac.il/

ORCID: 0000-0003-2332-5783

Address:
Joseph Keshet
Room 720, Zisaple Bldg
Faculty of Electrical sand Computer Engineering
Technion - Israel Institute of Technology
Haifa 3200003, Israel

Academic Degrees

2002 - 2008	Ph.D. in Computer Science, The School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel. Advisor: Yoram Singer
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.
1991 - 1994	B.Sc. in Electrical Engineering, Cum laude, Department of Electrical Engineering – Systems, Tel-Aviv University, Tel-Aviv, Israel.

Academic Appointments

2022 -	Associate Professor, Faculty of Electrical and Computer Engineering, Technion
2017 - 2022	Associate Professor, Department of Computer Science, Bar-Ilan University
2013 - 2017	Senior Lecturer, Department of Computer Science, Bar-Ilan University
2009 - 2013	Research Assistant Professor, TTI-C and Department of Computer Science, University of Chicago
2008 - 2009	Postdoc Fellow, IDIAP Research Institue at École Polytechnique Fédérale de Lausanne (EPFL).

Professional Experience

2022 - 2025	Chief Scientist, aiOla, Herzliya, Israel
2020 - 2022	Amazon Scholar, Amazon Alexa, Amazon Boston, MA
2020 - 2021	Founding Venture Partner, Disruptive AI Venture Capital, Herzelia, Israel
1995 - 2002	Head of speech processing research section, The R&D Center, Elite unit (8200), Israel Defense Forces (IDF)
1996 - 1998	Senior signal processing engineer, CTI^2 Ltd. (acquired by AudioCodes), Israel.
1994 - 1995	Signal processing engineer, Efrat Future Technologic (became Verint Inc.) , Israel.

Research Interests

I am excited about human speech. Speech is one of the most trivial yet one of the most complex signals we know. It bears information that the speaker would like to convey as well as her identity and her emotional and medical state. My passion for understanding and quantifying speech drives my research interests. My research concerns both machine learning and the computational study of human speech and language. My work on speech and language concentrates on speech processing, automatic speech recognition, speaker recognition, automating laboratory phonology, and pathological speech. My research on machine learning focuses on core machine learning and deep learning algorithms, specifically, that capture the structure of complex tasks, such as automatic speech recognition. But also – how to make them reliable and trustworthy.

Teaching Experience

Advanced Topics in Deep Learning: Transformers, graduate, Electrical and Computer Engineering, Technion (from Fall 2023)

Deep Learning for Speech Signals, Electrical and Computer Engineering, Technion (from Fall 2023)

Deep Learning, Electrical and Computer Engineering, Technion (2023)

Advanced Topics in Deep Learning: Speech Signals, graduate, Electrical and Computer Engineering, Technion (2022)

Introduction to Machine learning, undergraduate and graduate, Dept. of Computer Science, Bar-Ilan University (2014, 2015, 2016, 2017, 2019, 2020, 2021)

Advance techniques in machine learning, graduate, Dept. of Computer Science, Bar-Ilan University (2013, 2014, 2017, 2020, 2021)

Speech processing and recognition, graduate, Dept. of Computer Science, Bar-Ilan University (2017, 2018, 2019, 2021)

Probability for Computer Science, undergraduate, Dept. of Computer Science, Bar-Ilan University (2014, 2015, 2016, 2017, 2018)

Public Professional Activities

Deputy Editor, IEEE Transactions on Audio, Speech, and Language Processing (2026-now)

Senior Area Editor, IEEE/ACM Transactions on Audio, Speech, and Language Processing (2025-2026)

Elected member of IEEE Speech and Language Processing Technical Committee (2023-2025)

Guest Editor, IEEE Open Journal of Signal Processing (2023-now)

Associate Editor, IEEE Signal Processing Letters (2023-2024)

Associate Editor, IEEE/ACM Transactions on Audio, Speech, and Language Processing (2023-2025)

Journal reviewer: 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 ICASSP, Senior Area Chair in International Joint Conference on Artificial Intelligence (IJCAI), Neural Information Processing Systems (NeurIPS), 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

Organizing the second Israeli seminar on AI for Speech and Audio Processing, iSpeech 2025 (with Yossi Adi)

Organizing the first Israeli seminar on Speech and Audio Processing using Neural Nets, iSpeech 2023 (with Yossi Adi and Tal Rosenwein)

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 Learning^2," 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.

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

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

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.

Organizer of Illinois Speech Day 2011.

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.

Membership in Professional Societies

IEEE, member since 2002, senior member since 2017

ISCA (International Speech Communication Association), member since 2002

Fellowships, Awards and Honors

2026-2027	ISCA Distinguished Lecturer
2017	IEEE Senior Memeber.
2013	Israeli Council for Higher Education (VATAT) Fellowship for Oustanding Scholars
2010	Founder and chair of the Machine learning special interest group (SIG)
	of International Speech Communication Society (ISCA).
2007	Ph.D. Label Award, European Union PASCAL Network of Excellence.
2004-2006	Leibniz Student Fellowship.
2002	IBM Fellowship.
2001	Israel Defense Forces (IDF) – National award for outstanding research and development achievements.

Grants

2026	Technion security grant: Signal Enhancement for AI-Based Intelligence Systems in Operational Military Environments
2025	Rambam-Technion AI grant: First Clinical Acoustic Biomarker in Myasthenia Gravis Patients with Bulbar and Velopharyngeal Dysfunction
2024–2025	PMRI – Peter Munk Research Institute - Technion: Enhanced Automatic Speech Recognition Robustness to Military Noisy Environments
2024–2028	NIH R01: A Family-Genetic Study of Autism and Fragile X Syndrome (subcontractor)
2024–2028	NIH R01: Motoric features of the voice as indicators of psychomotor agitation and retardation in current and remitted MDD (subcontractor)
2023–2027	NIH 2-R01-DC010191-12: A Family-Genetic Study of Language in Autism (subcontractor)
2023	Elbit grant no. 131843: Unsupervised Estimation of Speech Intelligibility
2022–2025	NSF-BSF-NCS-FO: Enhancing speech and deep learning research through holistic acoustic analysis, with Matt Goldrick, PI
2022–2023	Modern Hebrew: Transcribed Speech Corpus (with Yossi Adi, The Hebrew University), Kamin, Israel Innovation Authority (660,000 ILS)
2021–2022	Explainability of Speech Recognition and Speaker Verification Systems, Magnet TrustAI, Israel Innovation Authority (610,000 ILS)
2020	Wireless Networks 5G, Magnet WIN, Israel Innovation Authority (301,000 ILS)
2020	AI-based Speech Therapy (with Dr. Yaki Cohen, Rambam Medical Center), Kamin, Israel Innovation Authority (660,000 ILS)
2020–2022	Generating artificial examples to improve the training of supervised networks using GANs and VAEs, Magnet AVATAR, Israel Innovation Authority (777,000 ILS)
2019–2020	Using Speech Acoustics to Reveal Motor Disruptions, *subcontractor*, with M. Goldrick, PI, and V. Mittal, PI, National Institutes of Health (NIH), Research Grant R21 (54,000 USD)
2019	Bank Leumi, Research Grant (22,500 ILS)
2018	Audio biomarkers extraction from call center voice, Google Research Grant (20,000 USD)

2018	The Speech, Language and Deep Learning Lab, Canada Foundation, Research Grant (100,000 CAD)
2018–2019	Cyber-Security Center, The Prime Minister's Office, Research Grant
2017–2019	Adversarial examples: generation and defenses, NEC Japan research grant (1,000,000 ILS)
2016	Machine learning and deep learning computation on network modules, Mellanox Research Grant (77,500 ILS)
2015–2017	Anomaly detection, Magnet Infomedia, Israel Innovation Authority (160,000 ILS)
2015	Non-invasive ultrasonic assessment of left ventricular filling pressure (with GE Medical systems), MAGNETON GE, Chief Scientist, The Israeli Minister of Industry (721,000 ILS)
2013–2014	Phonetic echoes of cognitive disruptions in speech production (with M. Goldrick), National Institutes of Health (NIH), Research Grant R21 (330,000 USD)

Graduate Students

Completed PhD theses

Names	Years	Thesis	Current employment
Yossi Adi	2015-2019	On the Robustness of Deep Neural Models and Their Applications to Speech Science	Assistant Professor, The School of Computer Science & Engineering, The Hebrew University; Senior Research Scientist, Meta AI Research (formerly Facebook AI Research)
Tzeviya Fuchs Hoshen	2015-2022	Identifying Unseen Words	Postdoc Fellow, Miri Adler Lab, HUJI
Gaby Shalev	2018-2022	Improving the Robustness of Deep Learning Algorithms	Research Scientist, Amazon
Felix Kreuk	2018-2022	Applications of Deep Neural Models to Speech Processing	Research Scientist, Meta AI Research (formerly Facebook AI Research)
Yael Segal	2019-2025	Deep Learning Algorithms for Prediction and Localization Speech Objects	Research Scientist, aiOla

Completed MSc theses

Names	Years	Thesis	Current employment
Danny Karmon	2014-2015	Risk Minimization in Structured Prediction using Orbit Loss	Senior Applied Research Manager, Microsoft
Stav Buchsbaum	2014-2016	Multiclass Support Vector Machine with a Reject Option (co-advised with Zeev Zalevsky)	Algorithm Developer, IoTerra
Einat Naaman	2015-2017	Learning Similarity Functions for Pronunciation Variations	Applied Researcher, General Motors
Amir Gottlieb	2015-2017	Automatic Analysis of Doppler Echocardiography using Structured Prediction	Senior Data Scientist, Taboola
Shua Dissen	2014-2017	Formant Estimation and Tracking using Deep Networks	Speech Recognition Researcher, Gong.io; Cont. to PhD under my supervision

Yaniv Sheena	2016-2018	Speech Segmentation using Deep Structured Models	Machine Learning Engineer, Facebook
Talia Ben-Simon	2019-2020	Learning Algorithms for Correcting Children's Speech Production	Cont. to PhD under my supervision
Gal Lev Shalev	2020-2021	Redesigning the Classification Layer by Randomizing the Class Representation Vectors	Research Scientist, General Motors
Yona Coscas	2019-2023	Photorealistic Image and Video Generation from Semantics	Director, Head of Artificial Intelligence - Elbit
Yosi Shrem	2019-2025	Learning Algorithms for Foreign Accent Conversion	Machine Learning Engineer, Meta

PhD theses in progress

Names	Years	Thesis
Roni Chernyak	2020-	Improving Adversarial Robustness in Images and Audio
Shua Dissen	2020-	New Methods in Un-supervised Speech and Speaker Recognition
Eyal Cohen	2022-	New Methods for Speech Recognition
Roi Benita	2022-	Diffusion Models for Speech (co-advised with Prof. Michael Elad)
Tomer Krichli	2024-	Adapting a Non-Causal Pre-Trained Speech Recognition Model for Streaming Tasks

MSc theses in progress

Names	Years	Thesis
Rotem Rousso	2023-	Neural Alignment Speech
Or Bitton	2024-	Can artificial deep neural networks read like the human brain? (Co-advised with Prof. Tzipi Horowitz-Kraus)
Vered Dassa	2024-	Explainability of ASR models
Gilad Nurko	2025-	Sampling for classification diffusion models
May Nathan	2025 -	Steganography using Diffusion Models
Omer Idovnik	2025-	Diffusion models for speech dereverberation
Itai Allouche	2025-	Improving speech decoding using reinforcement learning

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

Yael Segal-Feldman, Ann R Bradlow, Matthew Goldrick, **Joseph Keshet**, Open-vocabulary keyword spotting with hyper-matched filters for small footprint devices, Computer Speech and Language, 101986, 2026.

Itai Allouche, Itay Asael, Rotem Rousso, Vered Dassa, Ann Bradlow, Seung-Eun Kim, Matthew Goldrick, **Joseph Keshet**, How Does a Deep Neural Network Look at Lexical Stress?, The Journal of the Acoustical Society of America, Vol. 159, Issue 2, pp. 1348-1358, 2026.

Yehoshua Dissen, Shiry Yonash, Israel Cohen, and **Joseph Keshet**, A Front-End Adaptation Network for Improving Speech Recognition Performance in Packet Loss and Noisy Environments, IEEE Transactions on Audio, Speech and Language Processing, Volume 33, pp. 2175-2188, 2025.

Yael Segal-Feldman, Kasia Hitczenko, Matthew Goldrick, Adam Buchwald, Angela Roberts, **Joseph Keshet**, Enhancing analysis of diadochokinetic speech using deep neural networks, Computer Speech and Language, 101715, Volume 90, 2025.

Seung-Eun Kim, Bronya R. Chernyak, **Joseph Keshet**, Matthew Goldrick, Ann R. Bradlow, Predicting relative intelligibility from inter-talker distances in a perceptual similarity space for speech, Psychonomic Bulletin and Review, 2025.

Bronya R. Chernyak, Ann R. Bradlow, **Joseph Keshet**, Matthew Goldrick, A perceptual similarity space for speech based on self-supervised speech representations, Journal of the Acoustical Society of America, 155, pp. 3915–3929, 2024.

Seung-Eun Kim, Bronya R. Chernyak, Olga Seleznova, **Joseph Keshet**, Matthew Goldrick, Ann R. Bradlow, Automatic recognition of second language speech-in-noise, Journal of the Acoustical Society of America Express Letters, 4 (2), 2024.

Kasia Hitczenko, Yael Segal, **Joseph Keshet**, Matthew Goldrick and Vijay A. Mittal, Speech Characteristics Yield Important Clues about Motor Function: Speech Variability in Individuals at Clinical High-Risk for Psychosis, Nature Schizophrenia, Vol. 9, no. 60, 2023.

Eyal Cohen, Felix Kreuk, **Joseph Keshet**, Speech Time-Scale Modification with GANs, IEEE Signal Processing Letters, Vol. 29, pp. 1067-1071, April 2022.

Yael Segal, May Arama-Chayoth, **Joseph Keshet**, Pitch Estimation by Multiple Octave Decoders, IEEE Signal Processing Letters, Vol. 28, pp. 1610-1614, July 2021.

Adar Paz, Eshkol Rafaeli, Eran Bar-Kalifa, Eva Gilboa-Schectman, 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

Neta Glazer, Yael Segal-Feldman, Hilit Segev, Aviv Shamsian, Asaf Buchnick, Gill Hetz, Ethan Fetaya, **Joseph Keshet**, Aviv Navon Beyond transcription: Mechanistic interpretability in ASR, AAAI Conference on Artificial Intelligence, Volume 40, Issue 44, pp. 37407-37416, 2026.

Roi Benita, Michael Elad, **Joseph Keshet**, Spectral Analysis of Diffusion Models with Application to Schedule Design, Neural Information and Processing Systems (NeurIPS), 2025.

Gil Ayache, Menachem Pirchi, Aviv Navon, Aviv Shamsian, Gill Hetz, **Joseph Keshet**, WhisperNER: Unified Open Named Entity and Speech Recognition, IEEE Automatic Speech Recognition and Understanding Workshop (ASRU), 2025

Sarah Ethridge, Joe Lau, Bronya R Chernyak, Robert Voigt, Matt Goldrick, **Joseph Keshet**, Molly Losh, Self-Supervised Speech Representations in a Pre-train Speech Model Represent Key Rapid Automatized Naming Variability in Autism, Society for Computation in Linguistics, 8(1), 2025.

Aviv Navon, Aviv Shamsian, Yael Segal-Feldman, Neta Glazer, Gil Hetz, **Joseph Keshet**, FlowTSE: Target Speaker Extraction with Flow Matching, The 26th Annual Conference of the International Speech Communication Association (Interspeech), 2025.

Yael Segal-Feldman, Aviv Shamsian, Aviv Navon, Gill Hetz, **Joseph Keshet**, Whisper in Medusa's Ear: Multi-head Efficient Decoding for Transformer-based ASR, The 50th IEEE International Conference in Acoustic, Speech and Signal Processing (ICASSP), 2025.

Yehoshua Dissen, Shiry Yonash, Israel Cohen, **Joseph Keshet**, Enhanced ASR Robustness to Packet Loss with a Front-End Adaptation Network, The 25th Annual Conference of the International Speech Communication Association (Interspeech), 2024.

Aviv Shamsian, Aviv Navon, Neta Glazer, Gill Hetz, **Joseph Keshet**, Keyword-Guided Adaptation of Automatic Speech Recognition, The 25th Annual Conference of the International Speech Communication Association (Interspeech), 2024.

Arnon Turetzky, Or Tal, Yael Segal-Feldman, Yehoshua Dissen, Ella Zeldes, Amit Roth, Eyal Cohen, Yosi Shrem, Bronya R. Chernyak, Olga Seleznova, **Joseph Keshet**, Yossi Adi, HebDB: a Weakly Supervised Dataset for Hebrew Speech Processing, The 25th Annual Conference of the International Speech Communication Association (Interspeech), 2024.

Rotem Rousso, **Joseph Keshet**, Eleanor Chodroff, Tradition or Innovation: A Comparison of Modern ASR Methods for Forced Alignment, The 25th Annual Conference of the International Speech Communication Association (Interspeech), 2024.

Roi Benita, Michael Elad, **Joseph Keshet**, DiffAR: Denoising Diffusion Autoregressive Model for Raw Speech Waveform Generation, 12th International Conference on Learning Representations (ICLR), 2024.

Aviv Navon, Aviv Shamsian, Neta Glazer, Gill Hetz, **Joseph Keshet**, Open Vocabulary Keyword-Spotting with Adaptive Instance Normalization, The 49th IEEE International Conference in Acoustic, Speech and Signal Processing (ICASSP), 2024.

Gabi Shalev, Gal-Lev Shalev, **Joseph Keshet**, A Baseline for Detecting Out-of-Distribution Examples in Image Captioning, ACM Multimedia 2022, Lisbon, Portugal.

Bronya Roni Chernyak, Talia Ben Simon, Yael Segal, Jeremy Steffman, Eleanor Chodroff, Jennifer Cole, **Joseph Keshet**, DeepFry: Identifying Vocal Fry Using Deep Neural Networks, The 23rd Annual Conference of the International Speech Communication Association (Interspeech), Incheon, Korea, 2022.

Yael Segal, Kasia Hitczenko, Matthew Goldrick, Adam Buchwald, Angela Roberts, **Joseph Keshet**, DDKtor: Automatic Diadochokinetic Speech Analysis, The 23rd Annual Conference of the International Speech Communication Association (Interspeech), Incheon, Korea, 2022.

Yosi Shrem, Felix Kreuk, **Joseph Keshet**, Formant Estimation and Tracking using Probabilistic Heat-Maps, The 23rd Annual Conference of the International Speech Communication Association (Interspeech), Incheon, Korea, 2022.

Yehoshua Dissen, Felix Kreuk, **Joseph Keshet**, Self-supervised Speaker Diarization, The 23rd Annual Conference of the International Speech Communication Association (Interspeech), Incheon, Korea, 2022.

Talia ben Simon, Felix Kreuk, Faten Awwad, Jacob T. Cohen, **Joseph Keshet**, Correcting Misproducted Speech using Spectrogram Inpainting, The 23rd Annual Conference of the International Speech Communication Association (Interspeech), Incheon, Korea, 2022.

Tzeviya Sylvia Fuchs, Yedid Hoshen, **Joseph Keshet**, Unsupervised Word Segmentation using K Nearest Neighbors, The 23rd Annual Conference of the International Speech Communication Association (Interspeech), Incheon, Korea, 2022.

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, AAAI/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 Hejná, 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 the 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 Predic-

tion, 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 Begio, 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.