

Efthymios Tzinis

CONTACT INFORMATION

Primary-Email: etzinis@gmail.com **University-Email:** etzinis2@illinois.edu
Website: www.etzinis.com **Scholar:** [profile](#) **GitHub:** [etzinis](#) **LinkedIn:** [etzinis](#)


RESEARCH INTERESTS

Audio-visual Signal Processing, Sound Source Separation, Unsupervised & Self-Supervised Deep Learning, On-device/Federated Learning

EDUCATION




University of Illinois at Urbana Champaign (UIUC) Aug 2018 - Present

PhD Candidate in Computer Science (CS)

- Area: Artificial Intelligence
- Current Cumulative GPA: *4.00/4.00*
- Selected Courses: *Computational Inference and Learning, Machine Learning for Signal Processing, Optimization in Computer Vision, Machine Learning, Deep Learning, Deep Learning Theory, Probabilistic and Approximate Computing*
- Supervisor: Prof. Paris Smaragdis 

National Technical University of Athens (NTUA) Oct 2012 - June 2018

Diploma (BS + MEng) in Electrical & Computer Engineering (ECE)

- Highest Honors (top 2%), Cumulative GPA: *9.36/10.00*
- Specialization: Computer Science, Major (2 last years) GPA: *9.56/10.00*
- Thesis: *Manifold Learning and Nonlinear Recurrence Dynamics for Speech Emotion Recognition on Various Timescales*  
- Supervisor: Prof. Alexandros Potamianos 

Arsakeio Tositseio Ekalis Lyceum, Athens, Greece Sept 2009 - July 2012


Apolytirion, Highest Honors

- Top 0.1% in national qualifying exams (score: 19,713/20,000), GPA: *19.5/20*

PROFESSIONAL EXPERIENCE



Facebook Inc., Redmond, Washington, USA May 2021 - Aug 2021

Research Intern at Facebook Research Labs (FRL) (Working remotely)

- Self-supervised source separation at scale
- Supervisor: Dr. Anurag Kumar 



Google LLC, Urbana, Illinois, USA Aug 2020 - May 2021



Student Researcher at Google AI Perception (Working remotely)



- In-the-wild audio-visual universal sound source separation of on-screen sounds
- Next generation of spatio-temporal alignment of universal on-screen sounds
- Supervisors: Dr. Scott Wisdom  & Dr. John Hershey 


Google LLC, Cambridge, Massachusetts, USA May 2020 - Aug 2020


Research Intern at Google AI Perception (Working remotely)


- Unsupervised single channel sound source separation
- State-of-the-art purely unsupervised performance with mixture invariant training
- Supervisors: Dr. Scott Wisdom  & Dr. John Hershey 

Google LLC, Cambridge, Massachusetts, USA May 2019 - Aug 2019
 Research Intern at Google AI Perception
 – Utilizing sound classification for improving universal source separation
 – Conditioning separation models using semantic representations of multiple sound classes
 – Supervisors: Dr. Scott Wisdom  & Dr. John Hershey 

Behavioral Signal Technologies, Los Angeles, California, USA May 2017 - July 2018
 Machine Learning Engineer (Working remotely)
 – Leading machine learning infrastructure framework development
 – Implementing graph-structured pipelines for model training and feature extraction
 – Developing efficient real-time speech processing APIs
 – Building cognitive-affect, ASR and speaker diarization models
 – Supervisors: Dr. Thodoris Giannakopoulos  & Dr. Nassos Katsamanis 

ATHENA Research Center, Marousi, Greece May 2016 - July 2018
 Research Assistant at Institute for Language and Speech Processing (ILSP)
 – European project *BabyRobot* supported by *Horizon 2020* grant
 – Real time speech emotion recognition and multimodal engagement detection
 – Supervisor: Prof. Alexandros Potamianos 

SBA Research, Technological University of Vienna, Austria July 2016 - Aug 2016
 Intern, International Student Exchange Association (IAESTE)
 – Project: Seatlock, automatic screen locker for increased computer security
 – Implementing circuit level connection for pressure sensors
 – Programming Bluetooth low energy microprocessor to transmit encrypted messages
 – Developing Linux and Windows applications for screen locking and monitoring
 – Supervisor: Dr. Adrian Dabrowski 

Ernst & Young, Athens, Greece July 2015 - October 2015
 Intern at IT Advisory department (external partnership with Piraeus Bank)
 – Piraeus Bank’s MySQL database maintenance
 – Financial data analysis and risk prediction
 – Supervisor: Evangelos Kaslis 

HONORS AND AWARDS

Google on the spot bonus (\$ 1,000) 2021
 Awarded a bonus for my 9-month part-time job as a student researcher considering the final outcome of the project and my code delivery efficiency.

Facebook’s PhD Fellowship 2021
 Finalist (top 3.5% of submitted applications)



















Apple’s PhD Fellowship 2020
 Nominated to represent UIUC to the national competition





























UIUC Computer Science Excellence Fellowship (\$ 6,740) 2018
 Awarded to recruit and support promising incoming CS graduate students

HiPEAC Student Challenge Certificate, Zagreb, Croatia 2017
 GPU parallelization of K-means algorithm by using sparse matrix operations (NTUA-team)






Joint 1st place in EESTech Machine Learning Challenge, Athens, Greece 2017
 Solving supervised classification problems (participation: 40 teams)

Participation & Distinction 2015-2017
 Programming contests: Code Jam, Hash Code, Codechef, IEEE Xtreme (top 5%)














	Paris Kanellakis Fellowship (\$ 2,336)	2014-2016
	Highest GPA in computer science department (NTUA-ECE) for 5th-8th semesters	
	“The Great Moment of Education” Eurobank EFG Scholarship (\$ 1,168)	2012
	Highest rank in national qualifying exams in my school (Score: 19,713/20,000)	
	Award from Cultural Society of Santorini	2012
	Entering 7th in National Technical University of Athens, ECE department	
	Distinction	2011
	Maths & Physics competitions for all Greek high schools	
UNIVERSITY SERVICE APPOINTMENTS	Research Assistant at the UIUC	
	Deploy low-cost sound separation models into-the-wild	Spring 2020
	Exploring deep priors for blind source separation	Fall 2019
	Deep mask inference for source separation	Spring 2019
	Unsupervised source separation using deep clustering	Fall 2018
	Teaching Assistant at the UIUC	
	CS 446 / ECE 449 Machine Learning	Spring 2021
	CS 598 Machine Learning for Signal Processing	Fall 2020
	Lab Assistant at the NTUA	
	Natural Language & Speech Processing	Fall 2017
	Pattern Recognition	Fall 2017
	Teaching Assistant at the NTUA	
	Signals & Systems	Fall 2017
OTHER TEACHING EXPERIENCE	Private Tutor, Athens, Greece	2013 - 2017
	Maths, Physics, Programming and Circuits	
	– Preparation of high school students for national qualification exams	
	– Volunteer for social educational school of Athens (students in financial need)	
	Algorithms, C programming, Differential Analysis	
	– Giving lectures and tutoring undergraduate students for their exams	
CONFERENCE PUBLICATIONS	C1 Tzinis, E. , Wisdom, S., Jensen, A., Hershey, S., Remez, T., Ellis, D. P., and Hershey, J. R., “Into the Wild with AudioScope: Unsupervised Audio-Visual Separation of On-Screen Sounds.” To appear in <i>International Conference on Learning Representations (ICLR)</i> , 2021.      	
	C2 Tzinis, E. [†] , Bralios, D. [†] , Smaragdis, P. “Unified Gradient Reweighting for Model Biasing with Applications to Source Separation.” To appear in <i>International Conference of Acoustics, Speech and Signal Processing (ICASSP)</i> , 2021.      	
	C3 Wisdom, S., Tzinis, E. , Erdogan, H., Weiss, R. J., Wilson, K., and Hershey, J. R., “Unsupervised Sound Separation Using Mixture Invariant Training.” In <i>Advances in Neural Information Processing Systems (NeurIPS)</i> , 2020.    Spotlight (top 4% of submitted papers)	
	C4 Pariente, M., Cornell, S., Cosentino, J., Sivasankaran, S., Tzinis, E. , Heitkaemper, J., Olvera, M., Stöter, F.R., Hu, M., Martín-Doñas, J.M. and Ditter, D., “Asteroid: the PyTorch-based audio source separation toolkit for researchers.” In <i>Proceedings of Interspeech</i> , 2020, pp. 2637-2641.   	



- C5 **Tzinis, E.**, Venkataramani, S., Wang, Z., Subakan, Y. C., and Smaragdis, P., “Two-Step Sound Source Separation: Training on Learned Latent Targets.” In Proceedings of *International Conference of Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 31-35.      **Oral Presentation**
- C6 **Tzinis, E.**, Wisdom, S., Hershey, J.R., Jansen, A. and Ellis, D.P., “Improving Universal Sound Separation Using Sound Classification.” In Proceedings of *International Conference of Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 96-100.      **Oral Presentation**
- C7 Venkataramani, S., **Tzinis, E.** and Smaragdis, P., “End-to-end Non-Negative Autoencoders for Sound Source Separation.” In Proceedings of *International Conference of Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 116-120.   **Oral Presentation**
- C8 **Tzinis, E.**, Venkataramani, S. and Smaragdis, P., “Unsupervised Deep Clustering for Source Separation: Direct Learning from Mixtures using Spatial Information.” In *Proceedings of International Conference of Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 81-85.     **Oral Presentation**
- C9 Paraskevopoulos, G., **Tzinis, E.**, Ellinas, N., Giannakopoulos, T. and Potamianos, A., “Unsupervised low-rank representations for speech emotion recognition.” In *Proceedings of Interspeech*, 2019, pp. 939-943.   
- C10 **Tzinis, E.** †, Paraskevopoulos, G.†, Baziotis, C., and Potamianos, A., “Integrating recurrence dynamics for speech emotion recognition.” In *Proceedings of Interspeech*, 2018, pp. 927-931.     **Oral Presentation**
- C11 **Tzinis, E.**, and Potamianos, A., “Segment-based speech emotion recognition using recurrent neural networks.” In *Proceedings of Affective Computing and Intelligent Interaction (ACII)*, San Antonio, TX, 2017, pp. 190-195.    **Oral Presentation**
- C12 Chorianopoulou, A., **Tzinis, E.**, Iosif, E., Papoulidi, A., Papailiou, C. and Potamianos, A., “Engagement detection for children with Autism Spectrum Disorder.” In Proceedings of *International Conference of Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5055-5059.   **Oral Presentation**

JOURNAL
PUBLICATIONS









- J1 **Tzinis, E.**, Wang, Z., Jiang, X., and Smaragdis, P., “Compute and memory efficient universal sound source separation.” In *Journal of Signal Processing Systems*, 2021 (to appear).  
- J2 Katsarou, M.-S.†, Karathanasopoulou, A.†, Andrianopoulou, A.†, Desiniotis, V., **Tzinis, E.**, Lagiou, M., Charmandari, E., Chrousos, G.-P., Drakoulis, N., “Frequency Distribution of β_1 -, β_2 -, β_3 - Adrenergic Receptors Genetic Variants in a Southeastern European Caucasian Population.” In *Frontiers in Genetics, IF: 4.151*.   

WORKSHOP
PUBLICATIONS

- W1 **Tzinis, E.**, Wisdom, S., Jensen, A., Hershey, S., Remez, T., Ellis, D. P., and Hershey, J. R., “Self-Supervised Audio-Visual Separation of On-Screen Sounds from Unlabeled Video.” In *NeurIPS SAS Workshop for Self-Supervised Learning for Speech and Audio Processing*, 2020.   
- W2 **Tzinis, E.**, Wang, Z., and Smaragdis, P., “Sudo rm -rf: Efficient Networks for Universal Audio Source Separation.” In Proceedings of *IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2020, pp. 1-6.      **Oral Presentation**
- W3 Wisdom, S., **Tzinis, E.**, Erdogan, H., Weiss, R. J., Wilson, K., and Hershey, J. R., “Unsupervised Speech Separation Using Mixtures of Mixtures.” In *ICML Workshop on Self-supervision in Audio and Speech*, 2020.   
- W4 Venkataramani, S., **Tzinis, E.** and Smaragdis, P., “A Style Transfer Approach to Source Separation.” In Proceedings of *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2019, pp. 170-174.  
















- W5 Wang, Z., Subakan, Y. C., **Tzinis, E.**, Smaragdis, P., and Charlin, L., “Continual Learning of New Sound Classes Using Generative Replay.” In Proceedings of *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2019, pp. 308-312.  




UNDER SUBMISSION
& PRE-PRINTS

- P1 **Tzinis, E.**, Wisdom, S., Remez, T., and Hershey, J. R., “Improving On-Screen Sound Separation for Open Domain Videos with Audio-Visual Self-Attention.” arXiv preprint arXiv:2106.09669 (2021).  
- P2 **Tzinis, E.**, Casebeer, J., Wang, Z., and Smaragdis, P., “Separate but Together: Unsupervised Federated Learning for Speech Enhancement from non-IID Data.” Submitted to *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)* 2021.  
- P3 Wang, Z., Casebeer, J., Clemmitt, A., **Tzinis, E.**, and Smaragdis, P., “Sound Event Detection with Adaptive Frequency Selection.” Submitted to *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)* 2021.
- P4 **Tzinis, E.**, “Bootstrapped Coordinate Search for Multidimensional Scaling.” *arXiv preprint arXiv:1902.01482* 2019.  
- P5 Paraskevopoulos, G. †, **Tzinis, E.** †, Vlatakis-Gkaragkounis, E.-V. and A. Potamianos, “Pattern search multidimensional scaling.” Submitted to *Journal of Machine Learning Research (JMLR)*, *IF: 2.450*, 2018.  

† The indicated authors contributed equally in each work

TALKS

- ***Improving On-Screen Sound Separation for Open Domain Videos with Audio-Visual Self-Attention.***
 - *Virtual Workshop* Sight and Sound CVPR, June 2021. 
- ***Into the Wild with AudioScope: Unsupervised Audio-Visual Separation of On-Screen Sounds***
 - *Virtual Conference* ICLR, May 2021.  
- ***Unified Gradient Reweighting for Model Biasing with Applications to Source Separation***
 - *Virtual Conference* ICASSP, June 2021.  
- ***Self-Supervised Audio-Visual Separation of In-the-wild On-Screen Sounds***
 - *Virtual* NeurIPS SAS Workshop, Dec 2020. 
 - *Virtual* Google AI, Perception Spotlight, Nov 2020. 
- ***Sudo rm -rf: Efficient Networks for Universal Audio Source Separation***
 - *Virtual* Google AI, Sense Reading Group, Nov 2020.
 - *Virtual Workshop* MLSP, Espoo, Finland, Sep 2020.  
- ***Two-Step Sound Source Separation: Training on Learned Latent Targets***
 - *Virtual Conference* ICASSP, Barcelona, Spain, May 2020.  
- ***Improving Universal Sound Separation Using Sound Classification***
 - *Virtual Conference* ICASSP, Barcelona, Spain, May 2020.  
- ***Unsupervised Deep Clustering for Source Separation: Direct Learning from Mixtures using Spatial Information***
 - *Conference* ICASSP, Brighton, UK, May 2019. 
- ***Manifold Learning and Nonlinear Recurrence Dynamics for Speech Emotion Recognition on Various Timescales***
 - *Thesis Defense* NTUA, Athens, Greece, June 2018. 

ACADEMIC SERVICE	Reviewer / Program Committee	
	<ul style="list-style-type: none"> ➤ Conferences: <ul style="list-style-type: none"> – Annual Conference on Neural Information Processing Systems (NeurIPS) – International Conference on Learning Representations (ICLR) – International Conference of Acoustics, Speech and Signal Processing (ICASSP) ➤ Journals: <ul style="list-style-type: none"> – Transactions on Audio, Speech and Language Processing (TASLP) ➤ Workshops: <ul style="list-style-type: none"> – Workshop on Self-supervised Learning for Speech and Audio Processing at NeurIPS – IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA) 	
STUDENT MENTORING	Dean Biskup (MSc UIUC)	Summer 2021 - Present
	– Federated learning for source separation	
	Dimitrios Bralios (BSc/MEng NTUA - Visiting Student UIUC)	Fall 2020 - Present
	– Model Biasing with Unified Gradient Reweighting	
	Xilin Jiang (BSc UIUC)	Fall 2020 - Spring 2021
	– Efficient Audio Source Separation	
	Zhongweiyang Xu (BSc UIUC)	Fall 2020
	– Self-Supervised Audio Source Separation	
	Sacha Jungerman (BSc UIUC)	Fall 2019
	– Graph-based Representations of Sounds	
PROGRAMMING SKILLS	Languages (Excellent):	Python, C, Matlab, Unix Bash and SQL
	Languages (Familiar with):	Java, C#, C++, ML and Assembly (80x86,AVR)
	Operating Systems:	Linux, MacOS, Windows
	ML Frameworks:	Pytorch, TensorFlow, Keras
	Agile Development:	Git, JIRA, Jenkins, Scrum, Kanban Board
LANGUAGES	Greek (Native), English (Fluent)	
AFFILIATIONS	IEEE - Member	2016 - Present
	IEEE Signal Processing Society - Member	2017 - Present
	Hellenic Student Association at the UIUC - President	2018 - 2019
OTHER INTERESTS	Soccer, Beach rackets, Guitar, Singing, Traveling	
REFERENCES AVAILABLE UPON REQUEST	Paris Smaragdis	Professor, UIUC 
	John Hershey	Research Scientist, Google 
	Scott Wisdom	Research Scientist, Google 
	Alexandros Potamianos	Associate Professor, NTUA 