EFTHYMIOS (THEMIS) TSAPRAZLIS

3710 McClintock Ave., RTH 318 ◇ Los Angeles CA, 90089, USA ◇ (213) 681-2106

tsaprazl@usc.edu ◇ timtsapras23.github.io ◇ Google Scholar ◇ LinkedIn

EDUCATION

University of Southern California Aug 2024 - Aug 2029 (expected)
PhD in Computer Science, advised by Prof. Shrikanth Narayanan *Los Angeles CA, USA*

National Technical University of Athens Sep 2018 - July 2024
BSc and MEng in Electrical and Computer Engineering *Athens, Greece*
Thesis: *Enhancing contrastive language-vision pre-training with generative dialogue*, supervised by Prof. Petros Maragos

RESEARCH EXPERIENCE

Signal Analysis and Interpretation Lab, USC Aug 2024 - present
Graduate Research Assistant *Los Angeles CA, USA*

- Develop context-aware privacy frameworks for multimodal AI and privacy-preserving learning methods optimized for scalable, efficient deployment, enabling reliable privacy-utility trade-offs in real-world systems.

University of Texas at Austin Nov 2022 - Aug 2024
Research Student *Austin TX, USA*
Performed research on contrastive learning for vision-language pretraining advised by Prof. Alex Dimakis.

- Adapted CLIP with synthetic Q&A augmentation and distribution-level objectives, achieving +6% accuracy gains and improving multimodal alignment and transferability.
- Extended CLIP with a BLIP-2–trained synthetic dialogue tower, using contrastive fusion to improve zero-shot retrieval while preserving its pretrained capabilities. (NeurIPS 2023 SSL Workshop)

Computer Vision and Signal Processing Group, NTUA
Robot Perception and Interaction Unit, Athena RC Jan 2024 - July 2024
Research Assistant *Athens, Greece*
Built a training-free exploration system combining SAM segmentation, CLIP embeddings, and adaptive visual memory for incremental entity classification and open-ended autonomy. (ICRA 2025)

NCSR "DEMOKRITOS" Sep 2022 - Nov 2022
Research Intern *Athens, Greece*
Developed an end-to-end framework for multimodal sentiment analysis.

PROFESSIONAL EXPERIENCE

Workable Software Oct 2023 - Dec 2023
Machine Learning Engineer Intern *Athens, Greece*
Collaborated with product engineers to integrate CV parsing into production, introducing segmentation strategies that boosted accuracy and improved automated hiring reliability.

RESEARCH PROJECTS

Alignment in Multimodal AI through Explainable Privacy Understanding *USC - Amazon* (Sep 2024 - present)

- Received gift award from the USC–Amazon Center on Secure and Trusted Machine Learning.
- Designed a legally grounded visual privacy taxonomy and benchmark that aligns AI evaluation with GDPR, HIPAA, and EU AI Act principles, exposing how vision–language models fail to reason about context-sensitive privacy risks and and supporting post-training alignment as privacy judges.
- Building a privacy advisory agent to detect, localize, and sanitize sensitive visual content, to fulfill legal compliance with deployable generative-AI solutions.

Context-Aware Anonymization *USC* (February 2025 - present)

- Develop contextual privacy auditors that quantify leakage in natural conversations by identifying hidden attributes and decomposing text or speech into parallel information flows, enabling more explainable privacy evaluation.
- Design selective anonymizers using LLMs to remove or abstract only contextually sensitive flows while preserving task utility, improving privacy–utility.
- Extend framework to healthcare and legal domains, aligning flow-level privacy judgments with regulatory standards and building benchmarks for context-dependent anonymization.

Privacy Trade-offs in Speech Processing *USC* (May 2025 - present)

- Reframed speaker and attribute verification as membership inference attacks, showing that models with similar EER can leak up to 100× more at low-FPR, where attackers make reliable predictions, revealing that standard metrics hide true privacy risks.
- Investigating trade-offs in speech anonymization across privacy, downstream utility, perceptual quality, and computational cost.

IARPA HIATUS *USC ISI - UMD - UMich - UBirmingham* (Aug 2025 - present)

- Define and lead the lab's research agenda on authorship privacy for a multi-institution IARPA program; coordinate technical efforts, drive project next-steps, and present findings directly to IARPA program managers and partners.
- Building a spoken language privacy benchmark for attribute prediction, evaluating LLMs as both predictors and anonymizers, measuring privacy leakage with DP-inspired metrics, and analyzing style shifts from original to generated text.

SKILLS

Programming Languages	**Python**, C/C++ MATLAB/SIMULINK, SQL
Tools	**PyTorch**, Opacus, Bash, Jupyter, Git, TensorFlow, Weights&Biases

AWARDS

1. **Gerondelis Foundation Scholarship, 2025**
   PhD scholarship for Greek students pursuing graduate studies in the U.S.
2. **USC–Amazon Center on Secure and Trusted Machine Learning Gift Award, 2025**
   Competitive seed grant supporting research on privacy-preserving and trustworthy AI.
3. **Thomaides Award for Scientific Publications, 2025**
   Recognized for outstanding research publications in 2023 during undergraduate studies.
4. **ARISE Stipend Award, 2021**
   Full tuition coverage for participation in the EIT Digital Summer School 2021.

PUBLICATIONS

- **E. Tsaprazlis**, T. Lertpetchpun, T. Feng, S. P. Karimireddy, S. Narayanan. *VoxGuard: Evaluating User and Attribute Privacy in Speech via Membership Inference Attacks.* arXiv preprint arXiv:2509.18413, 2025. *(accepted ICASSP 2026)*
- **E. Tsaprazlis**, T. Feng, A. Ramakrishna, R. Gupta, S. Narayanan. *Assessing Visual Privacy Risks in Multimodal AI: A Taxonomy-Grounded Evaluation of Vision-Language Models.* arXiv preprint arXiv:2509.23827, 2025.
- P. P. Filntisis, **E. Tsaprazlis**, P. Oikonomou, F. Mattioli, V. G. Santucci, G. Retsinas, P. Maragos. *Towards Open-ended Robotic Exploration using Vision-Inspired Similarity and Foundation Models. IEEE ICRA*, 2025.
- **E. Tsaprazlis**, G. Smyrnis, A. Dimakis, P. Maragos. *Enhancing CLIP with a Third Modality. NeurIPS Workshop on Self-Supervised Learning*, 2023.