

EFTHYMIOS (THEMIS) TSAPRAZLIS

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

tsaprazl@usc.edu ◊ [timtsapras23.github.io](https://github.com/timtsapras23) ◊ [Google Scholar](#) ◊ [LinkedIn](#)

EDUCATION

University of Southern California

PhD in Computer Science, advised by Prof. [Shrikanth Narayanan](#)

Aug 2024 - Aug 2029 (expected)

Los Angeles CA, USA

National Technical University of Athens

BSc and MEng in Electrical and Computer Engineering

Sep 2018 - July 2024

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

Graduate Research Assistant

Aug 2024 - present

Los Angeles CA, USA

- Research on *Visual Reasoning* and *Multimodal Privacy Preserving Machine Learning*.

University of Texas at Austin

Research Student

Nov 2022 - Aug 2024

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. ([SSL@NeurIPS 2023](#))

Computer Vision and Signal Processing Group, NTUA

Robot Perception and Interaction Unit, Athena RC

Jan 2024 - July 2024

Athens, Greece

Research Assistant

- Built a training-free perception system for open-ended robotic exploration, enabling incremental scene understanding without predefined categories. ([ICRA 2025](#))

PROFESSIONAL EXPERIENCE

Workable Software

Machine Learning Engineer Intern

Oct 2023 - Dec 2023

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*

- Developed a legally-grounded visual privacy taxonomy (GDPR/HIPAA/CCPA), establishing guidelines for visual privacy and identifying critical limitations in existing datasets and VLM evaluation. ([SAIAD@CVPR 2026](#))
- Proposed CPRT, a compositional privacy risk framework with a 10.4k-sample benchmark and fine-tuned VLMs (4B/8B) that match or exceed frontier model performance on visual privacy reasoning. ([ECCV 2026](#))
- Building an end-to-end privacy advisory system that uses the taxonomy to detect and prioritize sensitive elements, localize them, and sanitize images through generative editing.

Speech Privacy *IARPA HIATUS & ARTS - USC*

- Proposed a DP-inspired MIA evaluation framework for speech anonymization, introducing low-FPR regime analysis and novel attribute privacy definitions, showing EER systematically underestimates real-world privacy risk. ([ICASSP 2026](#))
- Proposed rigorous mathematical privacy definitions for speech with provable upper bounds and a differentially private speaker anonymization method in the wild via calibrated noise on speaker embeddings. (Re-Data@COLM 2026)
- Built a topic-controlled spoken language privacy benchmark over 2,700 speakers, showing LLMs represent a qualitatively stronger threat than embedding models, recovering 1 in 4 speakers at 1% FPR under cross-topic conditions. (*Under Review*)

Visual Reasoning *USC*

- Investigating how to ground visual reasoning in human attention by constructing eye-tracking-based reasoning paths on VQA examples, with the goal of training VLMs that explicitly incorporate the visual modality during reasoning.
- Exploring test-time scaling for visual inputs, studying how grounding reasoning steps in spatial and visual representations can improve reasoning quality and reduce computational cost at inference.

TALKS

1. 8th Workshop on Safe Artificial Intelligence for All Domains @ CVPR 2026

Talk: *Assessing Visual Privacy Risks in Multimodal AI (Oral)*

June 2026, Denver, CO

2. 4th Annual Symposium of USC-Amazon Center

Talk: *Assessing Visual Privacy with VLMs*

May 2025, Los Angeles, CA

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.

SELECTED PUBLICATIONS

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