

Ioannis (Yiannis) Tsiamas



AI Research Scientist with a PhD in Speech Translation and experience at Meta FAIR, Apple AI/ML, and Dolby. Core contributor to OmniSONAR and Omnilingual MT at Meta — two of the most massively multilingual projects to date, supporting thousands of languages.

My work focuses on scaling multilinguality to thousands of languages, with emphasis on cross-lingual sentence embeddings, teacher-student distillation for language and modality expansion, zero-shot speech translation via cross-modal alignment, and automatic translation quality estimation.

CONTACT INFORMATION

Barcelona, Spain 📍
ioannis.tsiamas@gmail.com ✉️
www.ytsiamas.com 🌐

x.com/JohnTsiamas 🐦
github.com/johntsi 🐙
linkedin.com/in/i-tsiamas 🔗
huggingface.co/johntsi 🗣️
scholar.google.com/IoannisTsiamas 🎓

EXPERIENCE

∞ Meta (FAIR)

AI Research Scientist — Germany (Remote) Sep 2025 - Nov 2025

- Short-term researcher at the Omnilingual Team to finalize the OmniSONAR and OMT projects.

AI Research Scientist Intern — Paris, France Mar 2025 - Aug 2025

- Part of the Omnilingual Team at FAIR
- Core contributor to the [OmniSONAR](#) and [Omnilingual MT](#) projects.
- **OmniSONAR.** Led the extension of the sentence embedding space from 200 to 4,000+ language varieties via a contrastive-based teacher-student distillation framework with 15× cross-lingual error rate reduction on BIBLE (1,560 langs) and 2× on FLORES (200 langs) and surpassing 70B LLMs by 15 chrF++ in 1,560→English translation.
- **BLASER 3.** Created the most multilingual reference-free MT quality estimation model to date, outperforming previous metrics like MetricX-24 and xCOMET in correlation with human judgments.
- **Massively Multilingual Data Curation.** Engineered data preparation and filtering pipelines for 4,000 languages, supporting both OmniSONAR and Omnilingual MT training.
- **Evaluation Datasets.** Contributed to the creation of novel multilingual and multi-centric evaluation datasets: Met-BOUQuET and BOUQuET ([ACL 2025](#)).

AI Research Scientist Intern — Paris, France Aug 2024 - Feb 2025

- Proposed a character-level sentence encoder obtained via teacher-student distillation, that maximizes knowledge transfer across languages, surpassing traditional subword-based models in text understanding tasks.
- The speech-enabled adaptation of the encoder achieves state-of-the-art results in speech translation, surpassing large supervised models like SeamlessM4T, despite being trained only with minimal supervision from ASR data ([ACL 2025](#)).

🍏 Apple (AI/ML) — Aachen, Germany

AI Research Scientist Intern Apr 2024 - July 2024

- Introduced ContraProST, an evaluation benchmark and methodology for prosody awareness in speech translation, using LLMs and controllable TTS to generate contrastive examples — revealing that current end-to-end and production models underutilize prosodic cues despite having direct access to the speech signal ([WMT 2024](#)).

 **Dolby** — Barcelona, Spain

AI Research Scientist Intern

Nov 2023 - Feb 2024

- Proposed a sequential contrastive framework for audio-visual learning that contrasts non-aggregated representations using multidimensional sequential distances — achieving up to 3.5× relative improvement in retrieval recall over previous methods ([ICASSP 2025](#)).
 - Contributed to a masked generative video-to-audio model combining a full-band audio codec with sequence-to-sequence parallel generation, achieving both high audio quality and temporal synchronicity with visual actions ([ECCV 2024](#)).
-



Universitat Politècnica de Catalunya — Barcelona, Spain

PhD Researcher

Mar 2021 - May 2026

- Proposed SHAS, a novel automatic audio segmentation method that was the first to enable E2E speech translation models to retain almost perfectly the manual segmentation quality, which since has become a gold standard segmentation method ([Interspeech 2022](#)).
- Introduced an efficient data augmentation methodology to address the data scarcity issue in speech translation by utilizing SHAS to re-segment the training data, yielding consistent gains of up to +5 BLEU points across several benchmarks, and especially in low-resource settings ([EMNLP 2023](#)).
- Proposed a Perceiver-based speech encoder with Dynamic Latent Access that matches Transformer baselines at a fraction of the complexity, with flexible deployment across computational budgets ([ICASSP 2023](#)).
- Developed a zero-shot speech translation method using CTC compression and Optimal Transport to align speech encoders with multilingual MT models without any paired speech translation data, surpassing even large supervised models, and reaching state-of-the-art results ([ACL 2024](#)).

Academic Reviewer

- Acted as reviewer for conference submissions for ICASSP, Interspeech, and ACL-affiliated venues.

Research Engineer

- Developed an AI-based DJ for collaboratively playing with a human DJ (back-to-back) by utilizing sentiment-based song embeddings. The project was showcased in live performances on the SONAR Music Festival and the Barcelona Mobile World Congress. ([Video](#))

Teaching Assistant

- Organized and taught the labs of the MSc course *Artificial Intelligence with Deep Learning*.

Thesis Supervisor

- Provided technical and theoretical supervisor on BSc and MSc students for their thesis projects.
-



Zeta Alpha — Amsterdam, The Netherlands

AI Research Scientist Intern

Nov 2019 - Jun 2020

- Built a pairwise passage ranking and abstractive question-answering model deployed in the first version of Zeta Alpha's AI-powered search engine.
-



VU University Amsterdam — Amsterdam, The Netherlands

Teaching Assistant

Sep 2018 - Jul 2020

- Taught the practical labs and graded assignments and exams.
- Courses: *Machine Learning for the Quantified-Self* and *Introduction to Data Science*

EDUCATION



Universitat Politècnica de Catalunya — Barcelona, Spain

Ph.D. in Artificial Intelligence

Robust and Data-Efficient End-to-End Speech Translation

Supervised by José A. R. Fonollosa and Marta R. Costa-jussà

Mar 2021 - May 2026
(expected defense date)



University of Amsterdam — Amsterdam, The Netherlands

M.Sc. in Artificial Intelligence — GPA: 8.3/10 (*cum laude*)

Complex Question Answering by Pairwise Passage Ranking and Answer Style Transfer [Pdf, Code]

Aug 2020



VU University Amsterdam — Amsterdam, The Netherlands

M.Sc. in Quantitative Finance — GPA: 7.7/10

An ensemble of Recurrent Neural Networks for High-Frequency Stock Market Classification

Thesis published in the Journal of Forecasting [Web]

Oct 2018



Athens University of Economics and Business — Athens, Greece

Diploma in Accounting and Finance — GPA: 7.1/10

Jun 2015

PUBLICATIONS

Technical Reports

Omnilingual SONAR: Cross-Lingual and Cross-Modal Sentence Embeddings Bridging Massively Multilingual Text and Speech [Web, Pdf] 2026
The Omnilingual SONAR Team (including [Ioannis Tsiamas](#) as core contributor)

Omnilingual MT: Machine Translation for 1,600 Languages [Web, Pdf] 2026
The Omnilingual MT Team (including [Ioannis Tsiamas](#) as core contributor)

Conferences

BOUQuET: dataset, Benchmark and Open initiative for Universal Quality Evaluation in Translation 2025

EMNLP 2025 [Web, Pdf, Dataset]
The Omnilingual MT Team (including [Ioannis Tsiamas](#))

Improving Language and Modality Transfer in Translation by Character-level Modeling 2025
ACL 2025 [Web, Pdf]
[Ioannis Tsiamas](#), David Dale, Marta R. Costa-jussà

Sequential Contrastive Audio-Visual Learning 2025
ICASSP 2025 [Web, Pdf]
[Ioannis Tsiamas](#), Santiago Pascual, Chungsin Yeh, Joan Serrà

Speech is More Than Words: Do Speech-to-Text Translation Systems Leverage Prosody? 2024
WMT 2024 [Web, Pdf, Code]
[Ioannis Tsiamas](#), Matthias Sperber, Andrew Finch, Sarthak Garg

Masked Generative Video-to-Audio Transformers with Enhanced Synchronicity 2024
ECCV 2024 [Web, Pdf, Demo]
Santiago Pascual, Chungsin Yeh, [Ioannis Tsiamas](#), Joan Serrà

Pushing the Limits of Zero-shot End-to-End Speech Translation 2024
ACL Findings 2024 [Web, Pdf, Code, Models]
[Ioannis Tsiamas](#), Gerard I. Gállego, José A. R. Fonollosa, Marta R. Costa-jussà

SegAugment: Maximizing the Utility of Speech Translation Data with Segmentation-based Augmentations 2023
EMNLP Findings 2023 [Web, Pdf, Code]
 Ioannis Tsiamas, José A. R. Fonollosa, Marta R. Costa-jussà

Explaining How Transformers Use Context to Build Predictions 2023
ACL 2023 [Web, Pdf, Code]
 Javier Ferrando, Gerard I. Gállego, Ioannis Tsiamas, Marta R. Costa-jussà

Efficient Speech Translation with Dynamic Latent Perceivers 2023
ICASSP 2023 [Web, Pdf, Code]
 Ioannis Tsiamas, Gerard I. Gállego, José A. R. Fonollosa, Marta R. Costa-jussà

SHAS: Approaching Optimal Segmentation for End-to-End Speech Translation 2022
Interspeech 2022 [Web, Pdf, Code]
 Ioannis Tsiamas, Gerard I. Gállego, José A. R. Fonollosa, Marta R. Costa-jussà

Workshops

Speech Translation with Foundation Models and Optimal Transport 2023
IWSLT 2023 [Web, Pdf]
 Ioannis Tsiamas, Gerard I. Gállego, José A. R. Fonollosa, Marta R. Costa-jussà

Pretrained Speech Encoders and Efficient Fine-tuning Methods for Speech Translation 2022
IWSLT 2022 [Web, Pdf, Code]
 Ioannis Tsiamas, Gerard I. Gállego, Carlos Escolano, José A. R. Fonollosa, Marta R. Costa-jussà

End-to-End Speech Translation with Pre-trained Models and Adapters 2021
IWSLT 2021 [Web, Pdf, Code]
 Gerard I. Gállego, Ioannis Tsiamas, Carlos Escolano, José A. R. Fonollosa, Marta R. Costa-jussà

The TALP-UPC Participation in WMT21 News Translation Task 2021
WMT 2021 [Web, Pdf]
 Carlos Escolano, Ioannis Tsiamas, Christine Basta, Javier Ferrando, Marta R. Costa-jussà, José A. R. Fonollosa

Journals

An Ensemble of LSTM Neural Networks for High-Frequency Stock Market Classification 2019
Journal of Forecasting [Pdf]
 Svetlana Borovkova, Ioannis Tsiamas

SKILLS

Infrastructure: FSDP, mixed-precision training, slurm, multi-node training, data pipelines at scale

Frameworks: PyTorch, fairseq2, HuggingFace, polars, LaTeX

Research Areas: Cross-lingual Representation Learning, Speech Translation, Machine Translation, Cross-modal Alignment, Contrastive Learning, Knowledge Distillation, Automatic Quality Estimation, Transformer Architectures, Self-supervised Learning, Multimodality

LANGUAGES

English, Greek, Spanish, German