CURRICULUM VITAE

# Gijs Wijngaard

Utrecht, The Netherlands **g@gijs.me** **gijs.me** **Google Scholar** **GitHub** **LinkedIn**

Research scientist with a PhD in machine learning for audio.
Building models and datasets that let machines listen to and reason about sound.

## Education

DECEMBER 2021 – DECEMBER 2025

**PhD Machine Learning in Audio-Language**
University of Maastricht · The Netherlands

Research focused on audio-language models for scenes and events, particularly advancing automated audio captioning and audio question answering by combining symbolic representations of audio with deep neural networks. Defended on 7 September 2026 (recording).

THESIS — Data and Semantic Reasoning in Audio-Language Learning

SUPERVISORS — prof. dr. Elia Formisano, prof. dr. Michel Dumontier

---

2020–2021

**MSc Information Science – Data Science**
University of Amsterdam · The Netherlands · *GPA – 7.5*

Studied a one year master in Data Science, where I learned machine learning algorithms and software skills for working with data in different domains.

THESIS — Unsupervised approaches for analysing COVID-19 conspiracies

SUPERVISORS — prof. dr. Antal van den Bosch, dr. Maarten Marx

---

2019–2019

**Minor Information Technology**
Monash University · Australia

Studied in Melbourne, Australia, where I took a minor in data science and data analytics.

---

2015–2020

**BSc Artificial Intelligence**
Utrecht University · The Netherlands

Finished my studies in Artificial Intelligence, with a focus on learning machine learning techniques.

THESIS — Clustering soccer players: investigating unsupervised learning on player positions

## Publications

Google Scholar: 65 citations, h-index 4 (October 2026)

### Peer-reviewed

2025 — **Wijngaard, G.**, Formisano, E., Esposito, M., & Dumontier, M. (2025). AudSemThinker: Enhancing Audio-Language Models through Reasoning over Semantics of Sound. In *Advances in Neural Information Processing Systems (NeurIPS 2025)*.

2025 — **Wijngaard, G.**, Formisano, E., Esposito, M., & Dumontier, M. (2025). Audio-Language Datasets of Scenes and Events: A Survey. In *IEEE Access* (vol. 13, pp. 20328–20360) doi: 10.1109/ACCESS.2025.3534621.

2023 **Wijngaard, G.**, Formisano, E., Giordano, B. L., & Dumontier, M. (2023). ACES: Evaluating Automated Audio Captioning Models on the Semantics of Sounds. In *2023 31st European Signal Processing Conference (EUSIPCO)* (pp. 770–774). doi: 10.23919/EUSIPCO58844.2023.10289793.

### Preprints

2025 **Wijngaard, G.**, Formisano, E., Dumontier, M., & Jitsev, J. (2025). AudioToolAgent: An Agentic Framework for Audio-Language Models. *arXiv preprint arXiv:2510.02995*.

2025 **Wijngaard, G.**, Formisano, E., Esposito, M., & Dumontier, M. (2025). Data-Balanced Curriculum Learning for Audio Question Answering. *arXiv preprint arXiv:2507.06815*.

## Professional Experience

JANUARY 2026 – PRESENT

**Research Assistant**
Jülich Supercomputing Centre, Forschungszentrum Jülich · Jülich, Germany
Researching AI scaling laws and audio-language models within LAION, with the goal of building a general purpose audio-language model. Part of the Scalable Learning and Multi-Purpose AI Lab (SLAMPAI).

MAY 2026 – SEPTEMBER 2026

**Research Internship**
Microsoft Research · Redmond, WA, USA
Summer research intern in the Audio and Acoustics Research Group, working on spatial audio-language models. Hosts: Sebastian Braun and Hannes Gamper.

APRIL 2021 – JULY 2021

**Research Internship**
Meertens Institute · Amsterdam
Researching NLP techniques for analysing diversity in language and culture in the Netherlands. Part of the BETTER-MODS project, funded by the Dutch Research Council (NWO).

JUNE 2017 – NOVEMBER 2018 & SUMMER 2020

**Front-end Developer**
Channable · Utrecht
Development of a SaaS platform for managing and optimizing your online marketing. Built features on the frontend of the platform in Ember.js and React.

SEPTEMBER 2018 – FEBRUARY 2020

**Co-Founder & CTO**
Duwtje · Rotterdam
Co-founded an AI news application called Duwtje. Participated in Dutch Journalism Fund accelerator. Responsible for app development and data engineering.

## Teaching & Supervision

2022–2025

**Advanced Natural Language Processing**
University of Maastricht · KEN4259
Teacher of the MSc course (2023–2025) and teaching assistant (2022).

2023

**Information Retrieval and Text Mining**
University of Maastricht · KEN4153
Teacher of the MSc course.

2022–2024 **Introduction to Data Science**
Maastricht Summer School · MSS2001
Teacher for three editions of the summer school course.

2022–2025 **Research Project Artificial Intelligence**
University of Maastricht · KEN4130 / KEN4131
Supervised MSc student research projects, across both semesters, over three academic years.

## Talks & Posters

MAY 2026 *AI for Audio*. Invited talk, SURF, Amsterdam.

DECEMBER 2025 *AudSemThinker: Enhancing Audio-Language Models through Reasoning over Semantics of Sound*. Poster, NeurIPS 2025, San Diego, USA.

SEPTEMBER 2023 *ACES: Evaluating Automated Audio Captioning Models on the Semantics of Sounds*. Oral presentation, EUSIPCO 2023, Helsinki, Finland.

## Academic Service

2026 Reviewer, NeurIPS Evaluations & Datasets Track; ICLR

## Awards & Achievements

2026 Google Gemini Pro Research Grant. Grant: $10K

2025 NeurIPS Scholar Award – Travel Grant

2022–2025 Three Dutch Research Council compute grants for my PhD research. Grants: €50K, €100K, €100K

2024 OpenAI Research Grant for my work on AudSemThinker. Grant: $1000

2022 DCASE Travel Grant

2021 2nd Place in Course-wide Kaggle Challenge for Image Classification (out of 180+ students)

2019 Kairos Society Fellow – Fellow of a global community of top students and global leaders who aim to solve the world's greatest challenges.

2018 Winner Xomnia Datathon

2018 AWS Cloud Partitioner Essentials Certified

2018 Dutch Journalism Fund (Stimuleringsfonds voor de Journalistiek). Grant: €75K.

2017–2022 T500 Member – Was member of The Next Web's list of the most ambitious entrepreneurs under the age of 26.

2017 Winner ADEhack Hackathon

## Software, Models & Datasets

**AudioToolAgent** 22 ★
Agentic framework for audio-language models. Enables autonomous tool use and multi-step reasoning for complex audio understanding tasks. Implements ReAct-style prompting and tool integration for audio processing workflows.

**audio-datasets** 17 ★

Companion to the IEEE Access survey: Python scripts for dataset processing, visualization tools, and methods for mitigating data overlap in audio ML research. Dataset metadata released on Hugging Face.

**AudSemThinker** 15 ★

Multimodal model for audio-language understanding, with large-scale dataset creation from YouTube subtitles and training through SFT and GRPO. Three model variants and the AudSem dataset released on Hugging Face.

**ACES** 8 ★

Audio Captioning Evaluation on Semantics of Sound. Open-source Python package (`pip install aces-metric`) implementing a novel evaluation metric for audio captioning systems using RoBERTa-based semantic sound classification. Metric models released on Hugging Face.

## Skills

| | |
|---|---|
| PROGRAMMING | Python, Bash, Javascript, HTML, CSS, Java, C#, Prolog, Haskell |
| FRAMEWORKS | PyTorch, Hugging Face Transformers, Accelerate, DeepSpeed, SGLang, vLLM, Unsloth, WebDataset, Weights & Biases, Pandas, Gradio, Faiss, Polars, NumPy, SciPy, React |
| TOOLS | Git, Slurm, AWS, GCP, Rest API, Docker, Kubernetes, Jupyter, Kafka, Spark, ElasticSearch, SQLite, MySQL, PostgreSQL |
| CONCEPTS | Multimodal AI, Sounds and Natural Language Processing, Software Development, Functional programming, Object-oriented programming, Machine Learning, Deep Learning, Image Processing, Cloud Computing, Neurosymbolic Learning |