

Gijs Wijngaard

Curriculum Vitae

Utrecht, The Netherlands

✉ gijs@gijs.me

📄 gijs.me

🌐 linkedin.com/in/gijswijngaard

🎓 Google Scholar

Education

- 2021–2025 **PhD Machine Learning in Audio-Language**, *University of Maastricht*, The Netherlands.
Research focuses 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.
- committee prof. dr. Elia Formisano, prof. dr. Michel Dumontier
- 2021–2021 **MSc Artificial Intelligence**, *University of Amsterdam*, The Netherlands.
Started a second master in Artificial Intelligence, but dropped out to pursue my PhD.
- 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
- committee 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

First Author Publications

- 2025 Wijngaard, G., Formisano, E., & Dumontier, M. (2025). AudioToolAgent: An Agentic Framework for Audio-Language Models. Submitted to ICASSP 2025, under review.
- 2025 Wijngaard, G., Formisano, E., Esposito, M., & Dumontier, M. (2025). Data-Balanced Curriculum Learning for Audio Question Answering. *arXiv preprint arXiv:2507.06815*. Submitted to ICASSP 2025, under review.

- 2025 Wijngaard, G., Formisano, E., Esposito, M., & Dumontier, M. (2025). AudSemThinker: Enhancing Audio-Language Models through Reasoning over Semantics of Sound. Accepted to NeurIPS 2025.
- 2024 Wijngaard, G., Formisano, E., Esposito, M., & Dumontier, M. (2024). 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.
- 2021 Wijngaard, G. *Unsupervised approaches for analysing COVID-19 conspiracies*. Master's thesis, University of Amsterdam.
- 2020 Wijngaard, G. *Clustering soccer players: investigating unsupervised learning on player positions*. Bachelor's thesis, Utrecht University.

Professional Experience

- April 2021 – **Research Internship**, MEERTENS INSTITUTE, Amsterdam.
July 2021 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 – **Front-end Developer**, CHANNABLE, Utrecht.
November Development of a SaaS platform for managing and optimizing your online marketing.
2018 & Built features on the frontend of the platform in Ember.js and React.
Summer 2020
- September **Co-Founder & CTO**, DUWTJE, Rotterdam.
2018 – Co-founded a AI news application called Duwtje. Participated in Dutch Journalism
February 2020 Fund accelerator. Responsible for app development and data engineering.

Awards & Achievements

- 2022-2025 Several Dutch Research Council grants for my PhD research. Grants: €50K, €100K, €100K
- 2024 OpenAI Research Grant for my work on AudSemThinker. Grant: \$1000
- 2022 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

Selected Software Releases

- audio-datasets 14 ★ GitHub repository for the survey paper on audio-language datasets. Includes Python scripts for dataset processing, visualization tools, and methods for mitigating data overlap in audio ML research. Companion to IEEE Access publication.
- 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.
- AudSemThinker 5 ★ Multimodal AI model for audio-language understanding. Features large-scale dataset creation from YouTube subtitles, advanced training techniques (SFT and GRPO), and interactive demo deployment. Multiple model variants available 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