Stephen Sinclair

Amerikalaan 577, Utrecht 3526VX, the Netherlands
+31 (0)6 57 60 89 42 stephen.sinclair@nonnegativ.com
https://sinclairs.gitlab.io

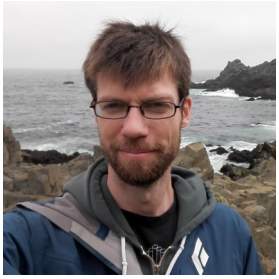

Professional Summary

Senior Machine Learning Engineer combining a robust software engineering background with a decade of AI-focused R&D in machine perception and action. I build deep learning solutions for multimodal data from computer vision to generative audio and text. My recent focus is on end-to-end development of commercial AI products, taking models from research to deployment as scalable, cloud-hosted services, with experience on AWS and Azure.

Skills

Programming Languages *(years)*

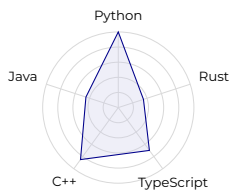

Technologies *(projects)*

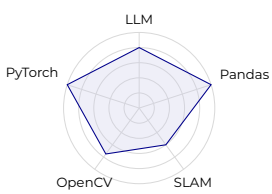

Topics *(projects)*

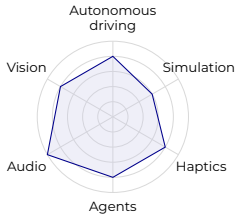

Tools & Methodologies

- ✔ AWS, Azure, Docker, k8s
- ✔ GitHub, GitLab
- ✔ PostgreSQL, DynamoDB
- ✔ MLFlow
- ✔ Agile, Scrum
- ✔ Jira, Notion

Spoken Languages

English: native
French: fluent
Spanish: fluent
Dutch: learning A2→B1

Employment History

Industry Positions

2022–2025 **Daisys BV**
Head of ML Research & Development, speech generation

2019–2022 **Navinfo Europe**
Senior Software Engineer, computer vision autonomous driving

2015–2018 **Inria Chile**
Software Engineer, robotics

2004–2005 **MPB Technologies Inc.**
Software Engineer, robotics

Research Positions

2018–2019 **Uni. Rey Juan Carlos**
Senior Research Fellow

2013–2015 **Institut Syst. Int. Robotiques (Paris 6)**
Research Fellow

Education

2005–2012 **McGill University**, (Montreal, Canada)
Master & PhD, audio-haptic robotics

2001–2004 **Concordia University**, (Montreal, Canada)
Bachelor's in Computer Science

Project Experience

Generative AI & LLMs

2025 **Agentic Pipeline, Language Learning** [link] *(Personal)*
(Solo) LangGraph agent for web data extraction and content synthesis, featuring "reflect-and-refine" loop to enforce quality thresholds; used DeepEval for quantitative performance validation. LangGraph | Structured decoding | DeepEval

2024 **Voice-interactive Agent** *(Daisys)*
Lead on a generative meditation app. LangGraph-driven user dialogue (voice-to-voice) gathers prefs for content gen; integrated Whisper and LLM with TTS streaming, AWS+Azure deployment. LangGraph | Whisper | OpenAI | Azure/AWS

2023 **Stylistic Adaptation & Generative Theatre** *(Daisys)*
Lead on a stylized content generation app using OpenAI for structured dialogue synthesis feeding into fine-tuned TTS for audio gen. Deployed on AWS (Lambda, DynamoDB), scalable GPU inference. Fine-tuning | LangChain | AWS

AI + Audio

2025 **Neural Speech Generation** *(Daisys)*
Hands-on lead for novel deep speech synthesis, from research (PyTorch models) to deployed API (FastAPI, AWS & Azure), and LLM demos. Novel solutions for natural, editable speech and voice. PyTorch | Transformer | NLP

2018 **Adversarial Autoencoder for Audio** *(Inria)*
Researched and implemented an AAE to clone the sound of physical models for real-time interactive audio synthesis. TensorFlow

2018 **Neural HRTF synthesis** *(Enosis + HP)*
Real-time neural inference written from scratch in C++ to predict HRTF filters for sound spatialization, embedded into a game engine and now driving an HP headphone product. C++ | Unity

AI + Vision

2022 **Few-shot Traffic Sign Classifier** *(Navinfo)*
Few-shot classifier for country-specific traffic signs based on spatial transformer and vector quantization of templates. ST+VQGAN | PyTorch | Classifier

2022 **Traffic Sign Localization** *(Navinfo)*
Developed a system for 3D localization of traffic signs from 2D object detection and semantic segmentation. OpenCV | ORB-SLAM | QGIS | YOLO

2021 **Highway Topology from Images** *(Navinfo)*
Designed and trained a constrained sequence model to infer highway lane connectivity and topology from image streams. LSTM+CRF | PyTorch

2021 **Highway Scenario Understanding** *(Navinfo)*
Built a pipeline to interpret and reconstruct complex driving scenarios from raw image data for model training and simulation. OpenCV | YOLO

2020 **Volume segmentation and deformation** *(URJC)*
Developed a cycle-consistent semantic segmentation model for analyzing optical coherence tomography images of human skin. Elastix | TensorFlow

Robotics, Haptics & Simulation

2017 **Physics Backend for Gazebo** *(Inria Chile)*
New physics backend for Gazebo open-source robotics simulator, improving accuracy for joint friction and contact constraints. C++ | Python

2012 **Haptic-Audio Interaction** *(McGill)*
Designed and optimized real-time physics and signal processing algorithms for multi-modal simulations. C++ | Embedded DSP | Python

Selected Publications

- **Sinclair, S. (2018).** "Sounderfeit: cloning a physical model using a conditional adversarial autoencoder." *Revista Música Hodie, 18.1, pp. 4460.*
- **Bochereau, S., Sinclair, S., & Hayward, V. (2018).** "Perceptual constancy in the reproduction of virtual tactile textures with surface displays." *ACM Transactions on Applied Perception (TAP), 15(2), 10.*
- **Pacchierotti, C., Sinclair, S., et al. (2017).** "Wearable haptic systems for the fingertip and the hand: taxonomy, review, and perspectives." *IEEE Transactions on Haptics.*
- **Sinclair, S. & Wanderley, M. M. (2008).** "A run-time programmable simulator to enable multimodal interaction with rigid-body systems." *Interacting with Computers, 21(1-2), pp. 54-63.*