Davide Bonura

davide.bonura.dev@gmail.com | linkedin.com/in/davide-bonura | github.com/dqvid3 | Palermo, Italy

Experience

Speech & Audio Processing Lab — University of Palermo Sep. 2025 – Mar. 2026
MSc Thesis | Supervisors: Prof. Sabato Marco Siniscalchi, Dr. M. La Quatra *Palermo, Italy*

- Designed a multimodal screening framework for Alzheimer's disease integrating raw handwriting sequences (CNN-Transformer + kinematic RF baseline) with ASR-generated speech transcriptions (Qwen3-ASR + fine-tuned UmBERTo); applied decision-level late fusion across two binary tasks (Healthy Control (HC) vs. Mild-AD and HC vs. Mild Cognitive Impairment (MCI)).
- Curated and standardised a novel dataset of 140 Italian-speaking participants; best system achieved macro-F_1 0.88 and clinical sensitivity 0.91 on HC vs. Mild-AD; handwriting kinematic analysis reached macro-F_1 0.75 on HC vs. MCI.

Speech & Audio Processing Lab — University of Palermo Mar. 2026 – Present
Research Collaborator | Supervisor: Prof. Sabato Marco Siniscalchi *Palermo, Italy*

- Extending the dataset with additional participants and preparing a peer-reviewed publication on the multimodal Alzheimer's screening system.

Education

University of Palermo Palermo, Italy
MSc in Computer Engineering, AI Curriculum — 110/110 cum laude (highest honours) *Mar. 2024 – Mar. 2026*

University of Palermo Palermo, Italy
BEng in Computer Engineering — 100/110 *Sep. 2020 – Mar. 2024*

Projects

Alzheimer's Speech Detection | *Wav2Vec2, WhisperX, BERT, ADReSSo-2021* 2024

- Re-implemented all 5 models from Pan et al. (INTERSPEECH 2021), replacing Kaldi ASR with Wav2Vec2 and WhisperX; implemented KenLM-guided hypothesis sweeping and beam-search diversity to reconstruct lattice uncertainty without raw lattice access.
- Reproduced acoustic-only, linguistic-only, fusion, and ASR-confidence-weighted classifiers on ADReSSo-2021; confidence-aware model achieved 85.92% accuracy, surpassing the original paper's 84.51%.

Speech Depression Assessment | *PyTorch, SVM, CNN, Wav2Vec2, WavLM, HuBERT, Whisper* 2025

- Compared SVM, CNN, and SSL-based classifiers for depression detection on the DAIC-WOZ and E-DAIC-WOZ datasets; conducted a systematic layer-level sweep across four SSL models (Wav2Vec2, WavLM, HuBERT, Whisper-small) to identify the most informative audio representations.

Tweet Irony & Emotion Classification | *PyTorch, BERTweet, HuggingFace, DeepSeek-V3, TweetEval* 2025

- Fine-tuned BERTweet-base on the TweetEval benchmark for irony detection (val F1: 0.767) and 4-class emotion classification (val F1: 0.815); applied domain-specific preprocessing (URL/mention normalisation, emoji-to-text, hashtag segmentation).
- Addressed class imbalance via LLM-based data augmentation: used DeepSeek-V3 to generate minority-class paraphrases; assessed semantic preservation using Twitter4SSE (domain-adapted tweet embeddings).

ASR-Kit | *Python, PyTorch, Qwen3-ASR* 2025 – Present

- Designed and built a modular Python ASR library with a unified `Transcriber` API and a pluggable driver registry; supports word-level timestamps via forced alignment and multi-file batch transcription.
- Distributed as a pip-installable package with optional per-backend dependencies; new model drivers can be registered without modifying core logic.

Multi-Object Pedestrian Tracker | *Python, PyTorch, OpenCV, DETR, MOT17* 2024

- Built a multi-object tracker using DETR for detection and Hungarian assignment; cost matrix combines IoU and cosine similarity on RoI-aligned CNN features. Evaluated on MOT17 (HOTA 25.8, MOTA 24.4 on test set).

Digital Signage Management System | *Java, JavaScript, Spring Boot, Jakarta EE, MySQL, Docker* 2024

- Built two independent Java microservices (Spring Boot 3.2 + Jakarta Servlet 6.0) for managing citywide LED billboard installations with real-time monitoring map (OpenLayers).
- Implemented role-based access control (Spring Security 6), XML schedule definitions (XSD/DTD schemas), geocoding autocomplete, and PDF/Excel report export (iText, Apache POI).

RubyRetail — Android E-Commerce App | *Kotlin, Android, Django, MySQL, Docker* 2023

- Built a full-stack e-commerce app: Kotlin/Material Design 3 frontend (Retrofit2 + OkHttp) consuming a Django REST API with MySQL; features include product browsing/search, cart, wishlist, checkout with saved addresses and payment cards, order history, and product reviews.

SPI Peripheral Design: RTL, FPGA & ARM32 | *SystemVerilog, C, ARM Assembly, Tang Nano 9K* 2025

- Implemented a custom 8-bit SPI counter peripheral across three abstraction levels: RTL in SystemVerilog with testbench, FPGA synthesis (Yosys/nextpnr on Tang Nano 9K), and a bit-banging bare-metal C driver on Raspberry Pi 3B+; extended with an ARM32 single-cycle processor simulation including custom SPI and GPIO peripherals.

Smart Warehouse Digital Twin | *C#, Unity 6, Neo4j, A*, Kalman Filter* 2025

- Developed an autonomous warehouse simulation with multi-forklift coordination: A* pathfinding, Kalman filter sensor fusion for positioning, and QR-code-based product sorting.
- Integrated Neo4j graph database for real-time state management; implemented an XAI layer broadcasting robot intent via on-screen UI.

Technical Skills

Programming: Python, Java, C, SQL, LaTeX, JavaScript, ARM Assembly, SystemVerilog, Kotlin, C#, MATLAB
Frameworks & Libraries: PyTorch, HuggingFace Transformers, scikit-learn, NumPy, Pandas, OpenCV, librosa, Spring Boot, Android SDK
Tools: Git, Linux, Jupyter, Docker, Unity, Neo4j, CometML
Languages: Italian (native), English (B2)

Interests

Weightlifting, Music, Chess, Cinema