# Daniele Re, PhD

ALGORITHM SCIENTIST & AI ENGINEER | SIGNAL PROCESSING, MACHINE LEARNING, LLM SYSTEMS, NEUROTECH
+972 50 920 7410 • +39 345 356 6416 • daniele.rejo@gmail.com • Givatayim, Israel
danielere.com • linkedin.com/in/daniele-re-phd-0a9664282 • Google Scholar • github.com/DanieleRe

## PROFILE

Scientist and engineer who turns weak, noisy signals from the body into measurements people can trust, and builds the software that puts them to use. PhD in cognitive neuroscience (Hebrew University), MSc in neuro-engineering and BSc in biomedical engineering (University of Genoa), followed by a postdoc in wearable neurotechnology (Tel Aviv University), an ML consultancy on dry-electrode wearables (X-trodes) and a clinical AI start-up I co-founded (Telephos), where I built and deployed the capture apps, the data platform and its agentic LLM layer myself.

Strongest at signal processing and sensor fusion, state-space and statistical modelling, deep learning on physiological time series, experimental design and validation against reference instruments. Hands-on across the stack: Python and PyTorch, FastAPI and React, Android/Wear OS, Docker and cloud GPU. 10 peer-reviewed publications, including Current Biology and Trends in Cognitive Sciences. Known for reporting null results and failure modes as carefully as wins.

## CORE COMPETENCIES

- Biosignal & sensor processing: EEG, EOG, EMG, ECG, PPG, EDA, IMU, eye tracking, voice and facial video; filtering, ICA, spectral and time-frequency analysis, drift correction, multi-stream synchronisation, automated signal-quality scoring.
- Machine learning: PyTorch transformers and sequence models, multimodal fusion, transfer from pretrained encoders, gradient-boosted trees and regularised models for small samples, anomaly detection, computer vision for faces and eyes.
- Estimation & statistics: Kalman filtering and state-space models, hierarchical Bayesian models (Stan), mixed-effects models, power analysis, multiplicity control, bootstrap, calibration, simulation-based identifiability, causal and trial methodology.
- AI engineering: LLM applications on the Claude API, tool calling and agent loops, structured outputs, RAG with citations, LangGraph workflows, voice agents, evaluation with gold sets and LLM judges, tracing, guardrails.
- Software & deployment: Python, TypeScript, Kotlin, C++, MATLAB, SQL; FastAPI, React, Next.js, Postgres, Databricks; Docker, Railway, Vercel, AWS, Cloudflare R2, Modal serverless GPU, Slurm; pytest, CI, Git, Linux.
- Neurotechnology & experiments: psychophysics, closed-loop systems, non-invasive stimulation (tACS, temporal interference), SimNIBS electric-field modelling, human-study design, protocol and safety documentation.

## PROFESSIONAL EXPERIENCE

### CTO / CSO & Co-founder | Telephos Intelligence (clinical AI start-up) 2025–present

Telephos measures behavioural and physiological state in dementia care, fusing facial video, voice and smartwatch physiology. I set the research direction and built most of the technical system myself.

#### Platform & AI engineering

- Built and deployed the platform (service.teleph0s.com) end to end: FastAPI backend with 17 REST endpoints, React front end, JWT authentication with role-based access, hash-chained audit logging, Postgres; Docker, auto-deploy on Railway, media on Cloudflare R2 with short-lived presigned URLs; 549 pytest tests and a daily data-quality watchdog.
- Shared LLM layer on the Anthropic API: Pydantic-validated structured outputs with validate–repair–retry, versioned prompts with rollback, PHI scanning, circuit breakers, cost and latency tracking, and end-to-end tracing (Arize Phoenix).

- Agentic features: a tool-calling session assistant that fetches only the time-series windows it needs and streams its reasoning to the UI (SSE); a patient-record assistant whose findings must cite source rows; natural language to charts over a closed catalogue; few-shot mapping of vendor files to a canonical schema.
- Designed a six-node LangGraph workflow for clinical-trial data quality: rule checks, root-cause clustering, LLM-generated hypotheses, draft site queries, human-review interrupts and restart-safe Postgres checkpoints.
- Evaluation before shipping: gold sets scored on finding recall, citation validity and calibration, LLM judges checked against human grading, regression suites for prompts.

**Voice agent & real-time avatar**

- Built a deployed voice check-in agent: ElevenLabs speech in and out, routed to our own state machine; questionnaire items read word for word, Claude writes only the conversational lines and each passes a rule-based guard (no advice, diagnosis, scores or numbers); crisis handling by fixed rules; memory across sessions, a personal baseline, English and Italian.
- Tested it with a bench of simulated patient personas and an LLM judge with never-cross checks, which caught missed crisis phrases and turn-handling bugs before real use.
- Prototyped a real-time video avatar driven by that agent: talking face rendered on a Modal GPU (Ditto, TensorRT, 46 fps), LiveKit WebRTC, browser face tracking (MediaPipe) so camera video never leaves the device, a versioned unit-tested reaction policy (nod, smile, concern, mirroring); about 1.3 s avatar delay, head motion benchmarked against Tavus; runs on phones.

**Capture systems & data pipelines**

- Specified the sensors and built the Android and Wear OS capture apps (Kotlin, Camera2, BLE) for clinic, home and caregiver use: facial video, voice, and a 26-channel continuous smartwatch stream (heart rate, EDA, skin temperature, motion).
- Built the watch-to-cloud data path: a durable on-device outbox so no data is lost across Bluetooth gaps, process kills and reboots, and incremental hourly uploads that cut per-sync volume about tenfold (3.3 MB to 0.3 MB).
- Feature pipelines on Modal and AWS: face detection, landmarks and action units; Whisper large-v3 transcription, speaker diarization (pyannote), WavLM arousal and valence, Claude labelling of distress and confusion; Databricks analytics and clustering with a plain-language query layer.
- Trained an on-device wearer-speech detector for the watch app (YAMNet embeddings + logistic head, AUC 0.92) and benchmarked on-device ASR (whisper.cpp) on FLEURS: 8% word error rate in English, 21% in Italian.

**Modelling, science & validation**

- Ran a simulation-based identifiability analysis that showed a planned feature could not be separated from reporting bias; dropped it and redesigned the measurement around sensor channels no reporter touches.
- Showed that a clinical voice classifier reporting AUC 0.96 was detecting the recording hardware, not the patients: a negative control built only from signal-free frames reached AUC 0.99, traced to device fingerprints.
- State-space and hierarchical Bayesian models (Kalman smoothing, Stan) for latent clinical state from sparse, irregular multimodal data, validated mainly by synthetic-truth recovery, ablations and calibration checks; chose a deterministic mixture-collapse scheme over a far larger amortised network for a cohort of a few hundred.
- Multimodal models on voice, video and clinical scales (openSMILE, wav2vec2 embeddings, transcripts) with leakage-free, subject-level validation; reported that an agitation classifier did not transfer to unseen residents.

**Clinical studies, regulation & communication**

- Wrote the protocol and launch plan for a UK clinic pilot (420 participants recorded) and data standards for ongoing US studies of agitation in dementia (dyads of patients and caregivers); estimands, prespecified analyses and covariate adjustment for trial methodology.

- Groundwork for regulated use: EU MDR/UKCA classification, GDPR/HIPAA data-flow mapping, de-identification, access control.
- Write the specifications engineers build from; present results and limits to clinicians, partners and investors; mentor engineers and scientists.

### ML Consultant (senior ML scientist scope) | X-trodes, Herzliya, Israel 2024–present

Consulting engagement on a dry-electrode wearable platform recording EEG, EOG, EMG, ECG and IMU outside the lab.

- Develop PyTorch transformer sequence models that recover eye movements from electro-oculography: a precise target reconstructed from a far weaker, noisier signal, prototyped as a hands-free control channel for head-worn devices.
- Validate against research-grade reference eye trackers, with baselines, ablations and subject-level protocols; map failure regimes under motion artefact, electrode noise, dropped samples and device-to-device variability.
- Build the pipelines that make recordings usable: filtering, artefact rejection and ICA, timestamp synchronisation across asynchronous streams, drift correction, automated signal-quality scoring, designed for low-latency inference.
- Hand over documented, reusable analysis code to the engineering team through Git, code review, automated tests and containerised environments.

### Postdoctoral Fellow | Tel Aviv University (Hanein lab), Israel 2023–2025

#### Closed-loop sensing & computer vision

- Built a real-time closed-loop system that detects movement events from wearable EMG and accelerometer streams and triggers an intervention within the movement window, validated with subject-level cross-validation.
- Built real-time camera pipelines for ocular behaviour: gaze and vergence (3D from 2D images), microsaccades, pupillometry and blink dynamics, integrated with electrophysiology.

#### Non-invasive retinal stimulation

- Led transorbital stimulation experiments (tACS, temporal interference) with facial dry-printed electrodes: rig set-up, stimulation protocol and safety design, stimulation-artefact handling in concurrent recordings.
- Built MRI-based, subject-specific planning with SimNIBS: head-model segmentation, finite-element electric-field simulation and montage optimisation (electrode placement and currents); the plans were used in real sessions.
- Showed that temporal interference tuned to the posterior retina increased phosphene detection in sighted participants, and that pulse timing on an oscillating envelope changes perception (first-author paper, IEEE EMBC 2025).

#### Experiments & collaboration

- Designed and analysed human studies end to end: protocols, power analysis, mixed-effects and hierarchical models, multiple-comparison control, psychometric fitting; mentored students; co-authored work on dry-printed electrodes for muscle stimulation and cognitive-load measurement.

### PhD Researcher | Hebrew University of Jerusalem (Landau lab), Israel 2017–2023

- Showed that attention samples visual information rhythmically, about 8 times per second for one object and 4 for two competing ones, across locations, features and the two eyes (Current Biology 2019; J. Cogn. Neurosci. 2023; TiCS 2026).
- Built dense behavioural time-course paradigms and the spectral and permutation statistics behind periodicity claims, and defended them in a published response to a methodological critique.
- Engineered a dichoptic setup on a 480 Hz PROPixx projector with circular polarisation to stimulate each eye separately; ran a flash-suppression binocular-rivalry study with continuous joystick report and EyeLink eye tracking.
- Wrote the experiment code (MATLAB, Psychtoolbox, DATAPixx/PROPixx control) and the behavioural and eye-tracking analyses; mentored students.

**Lecturer | University of Genoa, Italy** 2017–2018

- Taught natural scene statistics and neural coding to engineering students.

**Research Associate | King's College London, Department of Computer Science, UK** 2015–2016

- Implemented bio-inspired binocular receptive-field models in the iCub robot simulator: stereo disparity and depth from two calibrated views, driving vergence and saccadic camera control (first-author paper, VISAPP 2017).

## SELECTED PROJECTS

- **Agentic clinical-data assistant (Telephos).** Tool layer and agent loop so the model fetches only the data it needs; tool calls streamed to the UI so reviewers see the evidence. Checked with a gold set scored by an LLM judge, traces and unit tests.
- **Real-time AI video avatar (Telephos prototype).** Voice agent + GPU-rendered talking face over LiveKit with on-device perception and a tested reaction policy; latency and head motion measured against a commercial avatar.
- **Anomaly triage with a human in the loop (Telephos).** Automated checks flag anomalies in incoming trial data; flags are clustered, the LLM drafts a hypothesis per cluster as a structured output, and a reviewer accepts or rejects it, with every decision kept in the audit log.
- **Multimodal capture for a UK clinic pilot (Telephos).** Android and Wear OS apps (Kotlin, Camera2, BLE) recording facial video, voice and smartwatch physiology on a shared clock; feature-extraction pipelines on Modal and AWS; used in a pilot that recorded 420 participants.
- **Praxis (Next.js, TypeScript).** A practice app that grades free-text answers against rubrics with the Claude API and schedules review with spaced repetition.

## TECHNICAL SKILLS

- **Languages:** Python (primary), TypeScript/JavaScript, Kotlin, C++, C#, MATLAB, SQL, R, Stan, Bash.
- **Scientific Python:** NumPy, SciPy, pandas, statsmodels, scikit-learn, PyTorch, Hugging Face transformers, MNE-Python, MediaPipe.
- **Signal processing:** Filtering, ICA, spectral and time-frequency analysis, resampling, drift correction, event and onset detection, Lab Streaming Layer synchronisation, signal-quality metrics, artefact handling incl. stimulation artefacts.
- **Machine learning:** Transformers and sequence models, CNNs, multimodal fusion, transfer learning, gradient boosting, regularised regression, clustering (HDBSCAN), anomaly detection, dimensionality reduction; subject-level CV, leakage audits, calibration.
- **Statistics & modelling:** Kalman filtering and smoothing, switching state-space models, hierarchical Bayesian models, mixed-effects and longitudinal models, MMRM, power analysis, permutation tests, bootstrap, signal detection theory, psychometric fitting.
- **Speech, vision & sensing:** Whisper, faster-whisper, whisper.cpp, pyannote, WavLM, wav2vec2, openSMILE, YAMNet; face landmarks and action units; eye tracking (EyeLink), pupillometry; dry-electrode EEG/ExG, IMU, PPG, EDA; Android/Wear OS, BLE, Camera2.
- **LLM & agents:** Anthropic API, tool calling, structured outputs (Pydantic), RAG with citations, LangGraph, MCP, ElevenLabs, LiveKit, prompt versioning, LLM-as-judge evaluation, Arize Phoenix tracing, guardrails and PHI scanning.
- **Engineering & cloud:** FastAPI, React, Next.js, Postgres, Supabase, Databricks, Parquet; Docker, docker-compose, Railway, Vercel, AWS S3, Cloudflare R2, Modal serverless GPU, Slurm; pytest, CI, Git, Linux; agentic coding with Claude Code.
- **Neurotech & simulation:** SimNIBS (MRI segmentation, FEM electric fields, montage optimisation), tACS and temporal interference, Psychtoolbox, PsychoPy, PROPixx/DATAPixx, iCub simulator.

## EDUCATION

**PhD, Cognitive Neuroscience | Hebrew University of Jerusalem, Israel | 2017–2023**
Thesis: Attentional Sampling in Visual Hierarchy: exploring the dynamics of neural competition and perception (supervisor: Prof. Ayelet N. Landau).

**MSc, Neuro-Engineering | University of Genoa, Italy | 2014–2016**
Thesis: a bio-inspired model for binocular vision perception and oculomotor control (supervisors: Prof. Michael W. Spratling, King's College London, and Prof. Silvio P. Sabatini), carried out during a six-month research stay at King's College London.

**BSc, Biomedical Engineering | University of Genoa, Italy | 2010–2013**
Thesis: emergence of global oscillations in coupled neuronal networks (theory and numerical simulation; supervisor: Prof. Marco Storace). Project: spike-sorting algorithm for extracellular neural recordings (MATLAB).

## AWARDS, FELLOWSHIPS & PRIZES

**2025** InnoNext talent grant (€15,900), European Innovation Council, Horizon Europe (HORIZON-EIC-2023-TALENTS-01): funded innovation internship hosted by Altoida, October 2025 – January 2026.

**2019** JBC Doctoral Gold Fellowship ($15,000), Jerusalem Brain Community, Edmond & Lily Safra Center for Brain Sciences, Hebrew University.

**2019** Presidential Scholarship ($15,000), Jack, Joseph & Morton Mandel School for Advanced Studies in the Humanities, Hebrew University.

**2014** 2nd place, Italian Absolute Taekwondo Championship (black belt, July 2014; national and international competitor 2008–2016).

## TALKS, POSTERS & TEACHING

**2025** IEEE EMBC 2025 (presented in person): efficacy and usability of facial dry-printed electrodes for transocular stimulation.

**2018** ISCOP 2018, poster: Attentional sampling: the case of feature attention (Re & Landau).

**2017** VISAPP / VISIGRAPP 2017: binocular learned receptive fields for saccade-vergence control (paper presented).

**2017** Seminar, University of Genoa: From Natural Scene Statistics to Models of Neural Coding and Representation.

**2017–18** Lecturer, University of Genoa: natural scene statistics and neural coding.

## PUBLICATIONS

- **Re, D.**, Kusnir, F., & Landau, A. N. (2026). Attentional sampling resolves competition along the visual hierarchy. *Trends in Cognitive Sciences*, 30(2), 136–148.
- Teitelbaum, D., Ibrahim, R., Man, H., Luxembourg, N., **Re, D.**, & Hanein, Y. (2026). Multimodal quantification of cognitive load using a printed wearable facial bio-potential system. *Journal of Neural Engineering*, 23(4), 046029.
- **Re, D.**, Ibrahim, R., & Hanein, Y. (2025). Efficacy and usability of facial dry-printed electrodes for transocular stimulation. *IEEE EMBC 2025*, 1–6.
- Ibrahim, R., **Re, D.**, & Hanein, Y. (2025). Evaluation of dry printed electrodes for effective and safe transcutaneous electrical stimulation. *IEEE EMBC 2025*.
- Ibrahim, R., Vėbraitė, I., **Re, D.**, Funk, P. F., Volk, G. F., Guntinas-Lichius, O., & Hanein, Y. (2025). Dry-printed electrodes for transcutaneous electrical stimulation of innervated muscles: towards wearable and closed-loop stimulation. *J. NeuroEngineering and Rehabilitation*, 22, 226.
- Khasis, E., **Re, D.**, Kusnir, F., Bhat, A., & Landau, A. N. (2025). Sampling without vision: auditory performance fluctuates in congenitally blind individuals. *Journal of Cognitive Neuroscience*, 37(10), 1842–1854.

- **Re, D.**, Karvat, G., & Landau, A. N. (2023). Attentional sampling between eye channels. *Journal of Cognitive Neuroscience*, 35(8), 1350–1360.
- **Re, D.**, Tosato, T., Fries, P., & Landau, A. N. (2022). Perplexity about periodicity repeats perpetually: a response to Brookshire. *bioRxiv*.
- **Re, D.**, Inbar, M., Richter, C. G., & Landau, A. N. (2019). Feature-based attention samples stimuli rhythmically. *Current Biology*, 29(4), 693–699.
- **Re, D.**, Gibaldi, A., Sabatini, S. P., & Spratling, M. W. (2017). An integrated system based on binocular learned receptive fields for saccade-vergence on visually salient targets. *VISAPP*, 7, 204–215.

## LANGUAGES, SERVICE & INTERESTS

Languages: Italian (native); English (professional working proficiency).
Municipal councillor, San Lorenzo al Mare, Italy (2009–2014), alongside university studies.
Mentoring of engineers, scientists and students at Telephos and Tel Aviv University.
Interests: Taekwondo (black belt), running, travelling, cooking, music. References available on request.