Benan D. *Allan*

AI SYSTEMS ARCHITECT · PHD COGNITIVE & COMPUTER SCIENTIST

A decade of designing predictive intelligence systems for health and behavioural domains — interrogating what each data point truly represents, establishing evidential weight through statistical testing, and constructing variables when the signal doesn't exist in the raw data. Background spanning cognitive science and brain & computer science, applied to the foundational layer that sits above the modelling layer.

10+ yrs	6M	9 dom	PhD
PREDICTIVE SYSTEMS EXPERIENCE	USERS IN PRODUCTION · STARDUST	ON-DEVICE TEMPORAL CNN · NAR	PADOVA / EDINBURGH

CONTACT

benan@cogdev.ai

+33 7 88 51 52 08

Perpignan, France

ONLINE

benandemir.com

cogdev.ai · CID consultancy

Languages: EN · TR · IT · ES · FR

APPROACH

What does this data *actually mean?*

Most teams skip the foundational layer — what each data point truly represents, how much evidential weight it carries, whether the signal is even there to be modelled. **I start there.**

My work sits above the modelling layer: *representational data design*, evidential weighting through statistical testing, and constructing variables when the signal doesn't exist in the raw stream. The architecture decisions made before a model is trained — those are usually the ones that determine whether it works in production.

A decade across cognitive neuroscience, women's health, behavioural coaching, immersive tech, and quantitative finance. Different domains, same method.

CAPABILITIES

Six layers of an *intelligent system.*

From interrogating raw data to shipping the model behind a production endpoint.

01 / Architecture

Predictive system *architecture*

End-to-end design of the system that turns signal into prediction — data flow, feature stores, model registry, serving layer, feedback loops.

System design
Feature stores
Serving

02 / Representation

Representational *data design*

Deciding what each data point means, what it can carry, and how it should be encoded before it ever reaches a model. The decisions made here determine ceiling.

Schema
Encoding
Domain modelling

03 / Features

Feature *engineering*

Constructing variables when the raw signal doesn't carry them — from temperature-derived ovulation indicators to absorption ratios in orderbook microstructure.

BBT models
Microstructure
Derived signals

04 / Evidence

Evidential *weighting*

Statistical testing applied at the data layer — what should this point count for, how confident is the signal, when does the model defer instead of guess.

Hypothesis testing
Calibration
Uncertainty

05 / Behaviour

Cognitive & behavioural *modelling*

Modelling humans at the individual rather than aggregate level — personality, intention, trust, attention, decision-making. Built on a cognitive-science framework.

Personalisation
Profiling
Cognitive traits

06 / Deployment

Pipelines & *deployment*

From notebook to production endpoint — pipeline orchestration, VPS / serverless deploy, monitoring, and the boring infra that keeps inference running at 3 AM.

MLOps
Firebase
VPS / systemd

FEATURED PROJECTS

Production systems *I've shipped.*

Each build below shows the full stack — data layer, models, training methodology, serving infrastructure, and operational scale.

Stardust · Head of Science & AI · 2024 – 2026

PRODUCTION · 6M USERS

Stardust *Predictions*

BBT-BASED CYCLE, PERIOD & OVULATION PREDICTION AT 6M-USER SCALE

Built the AI function from the ground up — predictive models, data architecture, and ML infrastructure where none previously existed. BBT-derived models for cycle length, period length, and ovulation onset; spec-first feature design before training. Led a science team of 3 and guided an engineering team of 5 across research, architecture, and deployment.

6M	80%+
USERS IN PRODUCTION	REAL-TIME ACCURACY
90%+	3 + 5
OPTIMAL-USE ACCURACY	SCIENTISTS / ENGINEERS LED

MODELS

BBT trend detection · cycle-length regression · ovulation classifier

TRAINING DATA

BBT timeseries · cycle history · self-reported symptom & behavioural events

EVAL

80%+ real-time accuracy · 90%+ with optimal use, measured on next-cycle prediction

ARCHITECTURE

Data layer + feature store + ML infra + model registry — built ground-up

Python · Time series · BBT modelling · Feature engineering · Production ML · FemTech

SYSTEM NOTE

Inherited a product with no AI function. Designed and shipped: representational schema for the cycle-event data, BBT-derived feature constructors, a calibrated regression+classifier ensemble, a model registry, and the serving path back to the app — supporting 6M users in production.

CID · Women's health · 100% on-device ML

IN DEVELOPMENT

Nar

HORMONAL LIFE-STAGE INTELLIGENCE — CORE ML, ON DEVICE

iOS app covering 9 health domains across 6 hormonal life stages. Full science spec written before the model layer. All inference runs on-device — no server-side ML, no biometric data leaving the phone.

MODEL

Temporal CNN — 9 domain scores from 14-day biomarker windows (Float16, Core ML)

DATA INPUTS

HRV · resting HR · sleep stages (deep/REM) · SpO2 · respiratory · skin-temp Δ · steps · recovery — via HealthKit & Oura

ARCHITECTURE

NARKit internal SPM module · BaselineNormalizer · BiomarkerVector · LifeStageLens · TrendDetector

STACK

SwiftUI · Core ML (.mlpackage) · HealthKit · Oura SDK · NARKit · NarInsightsTestKit (XCTest)

Swift · Core ML · Temporal CNN · HealthKit · Oura · On-device · FemTech

CID · Infant nutrition

LIVE · APP STORE

TinyTaste

PERSONALISED INFANT NUTRITION TRACKING

iOS app for tracking solids introduction in 0–12 month old infants. Behind it: a nutrient profile database, DRI-based deficiency reasoning, and a recommendation engine that personalises against feeding history rather than aggregate norms.

DATA LAYER

Per-food `NutrientProfile` against 0–12mo DRI

REASONING

Deficiency detection · stage-aware suggestions

STACK

SwiftUI · Firebase Auth / Firestore · Apple Sign-In

PRICING

Freemium · $4.99/mo · $39.99/yr

Swift · Firestore · Nutrient modelling · iOS

CID · Behaviour change

BETA

AIM

AI-POWERED GOAL ACHIEVEMENT

An iOS goal-achievement app built on a cognitive-behavioural data model — goal decomposition, motivational state tracking, and AI-driven check-ins that adapt to the user's profile rather than running a fixed script.

DATA LAYER

Goal · subgoal · check-in · motivational state

MODELLING

Profile-conditioned LLM coaching · adaptive check-ins

STACK

SwiftUI · Firebase · Claude / GPT API

APPROACH

Behaviour-change theory grounded

Swift · LLM · Behaviour change · iOS

CID · Quant / Polymarket

LIVE · VPS

Sniperbot

MULTI-ASSET PREDICTION-MARKET TRADING

An autonomous trading system for Polymarket's price-touch markets across BTC, ETH, SOL, XRP. Five generations of architecture — from naive sniping to a modular family-state engine with shared blockers, regime ladders, and evidence-based promotion gates.

DATA LAYER

CLOB orderbook · whale wallets · structural microstructure

FEATURES

Touch probability · absorb ratio · net minting · sweep vs stealth

MODELLING

Family-state machine · per-family smart-sell · regime ladder

DEPLOY

VPS · systemd · Telegram alerts · auto-halt

Python · Polymarket CLOB · Microstructure · Systemd

PREDICTIVE SYSTEMS

Production AI *from the CV.*

Roles where I built or led the data and AI function — from architecture to deployment.

Stardust *App*

HEAD OF SCIENCE & AI

Jan 2024 – Feb 2026

Built the AI function from the ground up — designed and implemented all predictive models, data architecture, and ML infrastructure where none previously existed.

Developed **BBT-based prediction models** for cycle length, period length, and ovulation. Achieved 80%+ real-time accuracy and 90%+ with optimal use, serving **6 million users**.

Led a science team of 3 and guided an engineering team of 5 across research, architecture, and deployment.

6M	80%+	90%+
USERS SERVED	REAL-TIME ACCURACY	OPTIMAL-USE ACCURACY

Coach*Hub*

SENIOR DATA & BEHAVIOURAL SCIENTIST

2022 – 2023

Designed the **end-to-end personalisation framework from scratch** — mapping behavioural and personality data points across the full user journey to build dynamic user profiling.

Outcome: improved coach-coachee match quality by **86%** and lesson completion by **56%** within two weeks of deployment.

+86%	+56%	2 wks
MATCH QUALITY	LESSON COMPLETION	TO IMPACT

Geno*emote*

INTERIM CTO

Dec 2021 – 2022

Led research and technical strategy on **measuring and analysing brain activity as a physiological signal** for emotion recognition.

Designed the architecture of a trustworthiness rating system operating within VR environments — combining behavioural and neurological signals into a real-time assessment tool.

Humaine *AI*

LEAD COGNITIVE & COMPUTER SCIENTIST · CO-FOUNDER

Nov 2020 – 2023

Co-founded and led AI development of **five proprietary models** — including cognitive trait deduction from online behaviour, personality-driven product recommendation, and UX optimisation based on real-time user profiling.

Guided all model architecture and data design across client implementations, ensuring each system modelled human behaviour at an *individual* rather than aggregate level.

H-*Farm*

ADVANCED TECH LEAD RESEARCHER & PM

2018 – 2019

Led immersive technology development for **Bulgari, Gucci, Fendi, Adidas, and Ferrari** — VR shoe trial, virtual flagship stores, 3D product interaction systems.

Conducted behavioural data analysis of recorded in-VR human sessions to enhance system design, interaction patterns, and UX across all deployments.

CID

FOUNDER & CEO

March 2025 – Present

AI consultancy delivering **end-to-end solutions from architecture design to production deployment**. Specialise in context-aware AI systems for companies that need to understand human behaviour at a contextual level.

Design custom ML pipelines including data selection, model architecture, and fine-tuning strategies. Parent company for the projects above.

PIPELINE ANATOMY

From raw signal to *served prediction.*

A canonical pipeline I'll adapt to a domain — what changes is the signal source and the feature constructor; the spine stays the same.

01 · Source

Ingest

Raw events, sensor streams, orderbook, biometrics. Validate at the boundary.

02 · Mean

Representation

What does each point *mean*? Schema, encoding, missing-data semantics.

03 · Build

Features

Construct variables that don't exist in the raw stream — derived signals, aggregations.

04 · Weight

Evidence

Statistical testing — what should this point count for, with what confidence.

05 · Model

Predict

Model architecture, training, calibration, fine-tuning where it earns its keep.

06 · Serve

Deploy

Endpoint, monitoring, drift detection, feedback loop back to step 02.

FEATURE ENGINEERING

Variables that *weren't there yet.*

Derived signals built when the raw data didn't carry the answer — across domains.

BBT-derived ovulation STARDUST

Basal body temperature isn't a label — it's a noisy waveform. Built features that surface the thermal shift defining ovulation onset across irregular cycles.

Touch probability SNIPERBOT

Polymarket reach / dip markets resolve on path-dependent touches, not endpoints. 2× endpoint-probability fix replaces the naive mispricing estimate.

Absorb ratio & net minting SNIPERBOT V3

Structural microstructure signals built from anonymised orderbook flow — sweep vs stealth, absorption against the move, when whales are dark.

Cognitive trait deduction HUMAINE AI

Latent personality features inferred from clickstream and dwell behaviour — built to drive personalisation at the individual rather than segment level.

Match-quality features COACHHUB

Mapped behavioural and personality data points across the user journey to build a dynamic profile feeding match-ranking — +86% match quality.

Nutrient deficiency reasoning TINYTASTE

Per-food nutrient profile × 0–12mo DRI windowed against feeding history — features that drive next-food suggestion and gap detection.

VR-session behavioural fingerprint H-FARM

Recorded in-VR sessions reduced to behavioural features that drove iterative system, interaction, and UX improvements across luxury client deployments.

EEG / fMRI emotion features PHD · GENOEMOTE

Multimodal physiological-signal fusion (EEG, fMRI, GSR, eye-tracking) into emotion-recognition features for VR-based real-time assessment.

DEPLOYMENT

Where the systems *actually run.*

The unglamorous infra layer that keeps inference live.

VPS · systemd

Sniperbot lives here — patch script → scp → ssh → systemctl restart. Telegram alerts via shared bot. Auto-halt on drawdown.

Firebase

Auth, Firestore, Functions for TinyTaste & AIM. Apple Sign-In, Keychain persistence, server-side merge writes.

App *Store*

iOS distribution for TinyTaste & AIM. xcodegen-managed projects, version source-of-truth in `project.yml`.

GitHub *Pages*

Marketing surfaces — cogdev.ai, benandemir.com, this portfolio. Static, fast, custom domain via CNAME.

METHODS & TOOLS

From *brainwaves* to backends.

The kit underneath the work — research instruments, modelling stacks, and shipping infrastructure.

LANGUAGES

Python
data, ML, automation
Swift / SwiftUI
iOS production
JavaScript / Node
web, tooling
R
statistical research

ML / AI

scikit-learn · XGBoost
PyTorch · TensorFlow
Claude / GPT APIs
Fine-tuning · prompt engineering

DATA

pandas · numpy · scipy
Firestore · BigQuery · Postgres
Multimodal fusion
EEG · fMRI · GSR · eye-track
Orderbook / market microstructure

DEPLOY / OPS

Firebase · Cloud Functions
VPS · systemd · cron
GitHub · GitHub Pages · CI
Telegram alerts · monitoring

RESEARCH INSTRUMENTS

EEG
physiological signal
fMRI
neural correlates
GSR
autonomic arousal
Eye-tracking
attention & intent

STATISTICAL

Hypothesis testing
Mixed-effects models
Bayesian inference
Calibration · uncertainty quantification

DOMAINS

Women's health · FemTech
Behavioural science · coaching
Cognitive neuroscience · VR
Quantitative trading · prediction markets

FRAMEWORKS I USE

Representational data design
Evidential weighting
Spec-first ML
Individual-level modelling