CASE STUDY

Neuronia

Building a Mini On-Device LLM for Brain Training

Train. Play. Evolve.

ViRan Digital Labs | Singapore | March 2026
React Native | Expo SDK 54 | TypeScript 5.9

Executive Summary

An 8-year-old's game concept evolved into a 12-game cognitive training platform with 88,000+ questions, an adaptive ML engine, and 11 specialized scoring models — all running entirely on-device. Zero data collection. Zero cloud dependency. This case study documents the engineering decisions, adaptive intelligence architecture, and quality assurance framework that made it possible.

12	42+	88K+	361	0	100%
Games	Game Modes	Questions	Automated Tests	Data Collected	Offline

The Challenge

- **100% offline architecture —** every decision, every game runs without internet or server calls.
- **Adaptive ML engine —** learns player behavior, adjusts difficulty, prevents stagnation on-device.
- **88,000+ questions at scale —** generated via 25 parameterized templates across 49 categories.
- **12 distinct scoring engines —** each game has unique mechanics requiring specialized algorithms.
- **Multi-store compliance —** iOS + Android approval with GDPR, PDPA, CCPA, COPPA standards.
- **Zero data collection —** all intelligence built locally. No accounts, no tracking, no ads SDK.

KEY INSIGHT

"Could we build an intelligent, adaptive brain training engine — like a miniature LLM — running entirely on a phone?"

Origin

Game concept creator: Vizhinee Srinivassane (Age 8). What began as "Dad, can you make a game where you tap the right answers?" became Neuronia — a platform spanning 4 cognitive zones, 12 interconnected games, 42+ modes, and a composite scoring engine that builds a unified cognitive profile from every game played.

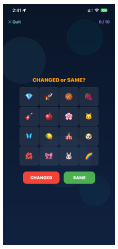

The Adaptive Intelligence Engine

Neuronia's core engine functions like a purpose-built language model running on-device. It maintains context, adjusts behavior from feedback, and makes probabilistic decisions — without a single network call. Below: how core ML/LLM concepts map to Neuronia's architecture.

Rolling Accuracy Window
Context Window

Tracks last 5 responses in a sliding window. Local context drives real-time difficulty adjustment — detecting improvement, plateau, or fatigue.

Difficulty Distribution
Temperature Control

Manages difficulty variance: >80% accuracy → 40% harder; <40% → 40% easier; 40-80% → balanced 60/20/20. Keeps players in optimal learning zone.

Anti-Stagnation Logic
Repetition Penalty

After 8+ questions without change, forces recalibration. Prevents cognitive equivalent of local optimum — ensures continuous growth.

Consecutive Tracker
Confidence Adjustment

3 correct → difficulty up; 2 wrong → difficulty down. Responds to streaks, not just averages — mirrors sequential feedback adjustment.

11 Scoring Engines
Domain-Specific Models

Each game has a specialized model: BrainRain combo multipliers, IQ accuracy-to-score mapping, Reaction RT/Accuracy/Consistency compositing, Speed Math 6-tier adaptive.

CognitivLab Ensemble
Model Aggregation

Zone-weighted ensemble: Brain Lab 30%, Reflex 25%, Knowledge 25%, Discovery 20%. Missing zones re-normalized. Multiple signals → robust assessment.

Question Generation Pipeline

25 parameterized templates generate 88,000+ unique questions with near-infinite replay value. All stored in optimized JSON — zero API dependency.

- 38,335 Quizzeria questions across 49 categories
- 11,785 IQ questions: logic, patterns, math, verbal, spatial
- 19,830 BrainRain elements across 3,917 groups
- 947 True/False statements for TF Blitz (5 modes)
- Deterministic daily challenges via seeded shuffle
- Procedural generation: 5K–40K variants per template

EQ Assessment

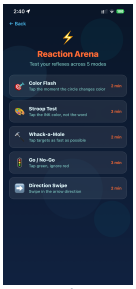

Reaction Arena

Scoring Engine Architecture

ENGINE	SCORING METHODOLOGY
BrainRain	+10 to +40 correct, combo multipliers 1.5–5.0x, accuracy bonus at 95%+
IQ Check	Accuracy-band IQ mapping: 90 at 20% → 140 at 90%+, speed bonuses, streak tracking
Reaction Arena	Composite = (RT×0.50 + Accuracy×0.30 + Consistency×0.20) × 2, Reflex Age mapping
Speed Math	6-tier adaptive: Tier 1 targets 1000ms, Tier 6 targets 3500ms, streak-based jumps
CognitivLab	Zone-weighted ensemble with auto re-normalization for missing data

Quality Assurance Framework

Cognitive assessment demands absolute reliability. A subtle scoring bug erodes user trust. Comprehensive testing infrastructure was built from day one.

361	29	100%	10	<2s
Test Cases	Test Suites	Pass Rate	Critical Bugs Caught	Full Suite Execution

Critical Issues Resolved

- NaN divisions in scoring engines
- Empty array crashes in profile aggregation
- Denominator errors in composite calculations
- Stale state in adaptive difficulty engine
- Memory leaks in long game sessions
- Race conditions in AsyncStorage persistence

Test Coverage Highlights

- 124 smoke tests for edge-case crash prevention
- Dedicated suites for all 11 scoring engines
- Adaptive difficulty engine: 13 test scenarios
- CognitivLab ensemble: 9 aggregation tests
- StorageManager: 42 persistence tests
- Full screen rendering test coverage

App Store Compliance

- **iOS App Store:** Full HIG compliance, first-submission readiness, zero third-party tracking SDKs in binary.
- **Google Play:** Complete Data Safety certification, no data collection/sharing, rated Everyone / PEGI 3.
- **Privacy:** GDPR, PDPA, CCPA, COPPA — all exceeded. Zero data = zero compliance risk.
- **Clean Binary:** No ads SDK, no analytics, no crash reporter. Only game logic + question DB + adaptive engine.

Key Recommendations

1. **Design for privacy at architecture level.** Zero-data-collection made compliance reviews trivial across 4 jurisdictions. Constraints breed elegance.
2. **Invest in generation pipelines over manual content.** 25 templates produced 88K+ questions. Scale comes from structure, not volume.
3. **On-device ML can replicate cloud intelligence.** Rolling windows, weighted distributions, and anti-stagnation create adaptive experiences without servers.
4. **Automated testing is non-negotiable for cognitive apps.** 10 critical bugs would have been 1-star reviews. A wrong IQ score destroys trust instantly.
5. **Composite scoring drives user engagement.** The unified CognitivLab profile motivates players to explore all 12 games to improve their overall score.

Neuronia — Train. Play. Evolve.

Game concept: Vizhinee Srinivassane (Age 8) | Engineered by: Srini Ramakichenane

ViRan Digital Labs | virandigitallabs.com | howsmartru.app@gmail.com | Singapore