

LAB NOTEBOOK · dated



THE LAB NOTEBOOK

A neural network on a 1982 ZX Spectrum

In April 2026 we ran a real neural text generator on a 44-year-old 8-bit machine: a 2.2 KB GRU on a Z80, multiplying by shift-and-add and doing sigmoid from a 256-entry lookup table. It proves neural inference needs neither a GPU nor floating point — and it's deliberately open: fakes don't run on a Z80.

The dated R&D archive

 Research — measured, partial

A neural network on a 1982 ZX Spectrum

In April 2026 we ran a real neural text generator on a 44-year-old 8-bit machine: a 2.2 KB GRU on a Z80, multiplying by shift-and-add and doing sigmoid from a 256-entry lookup table. It proves neural inference needs neither a GPU nor floating point — and it's deliberately open: fakes don't run on a Z80.



01 The entry, dated

02 What it was

03 What we built

04 What we learned — including the honest edges

05 Where it went / status

06 Sources, glossary & the honesty ledger

Every claim carries its tier · nothing faked

PREPARED

8 April 2026

VERIFIED BY

Jon Ross — Founder, Vocabotics —
15 years building safety-critical
systems

HOW TO READ IT

Drafted with AI, verified by a
named human. Sources at the back.

01 The entry, dated



The entry, dated

Apr 2026

entry date

Hardware

category

**the proof, not the
lead**

lead over the world



Public

access

What the world had at the time

The default assumption is that AI needs a GPU, floating-point maths, and gigabytes. A working neural net in 2.2 KB on a 1982 Z80 is a direct, physical refutation of that assumption — which is why it makes such a good open demo.

A real neural net, generating text, in 2,229 bytes on a computer older than most of the people reading this — no GPU, no floating point, no cloud.

The costly signal

The measured result

Every number carries its own honesty tier — the tiering is the credibility.

2.2 KB

the whole model, as a .tap file

 RESEARCH

Z80

1982 8-bit CPU, no FPU

 RESEARCH

shift+add

multiply-free integer maths

 RESEARCH

256-LUT

sigmoid from a lookup table

 RESEARCH

If you want to prove that AI does not *need* a datacentre, the cleanest possible demonstration is to run it on a machine from 1982. You cannot fake a result on a Z80 — it either fits in the memory and runs on the clock cycles, or it does not. This report is **open on purpose**: it is a viral, trust-building proof, and there is no moat to protect in it.

02 What it was



A project to run **neural inference on vintage hardware** — ZX Spectrum, BBC Micro, C64, Amiga, Psion, 486 — wrapped in a generative adventure game where the model's hidden state literally *is* the save file. The organising question: how low can the hardware floor for a real neural network go?

03 What we built



A real **2-layer INT8 GRU** exported to run on the metal:

A **2.2 KB .tap file** — the entire model — running on a **Z80**.

Multiply-by-shift-and-add integer maths, because the Z80 has no floating-point unit.

Sigmoid from a 256-entry lookup table, because you cannot afford to compute it.

 RESEARCH — MEASURED, PARTIAL

The measured fact is delightfully concrete: a **2,229-byte** neural text generator that runs on a **1982 8-bit CPU**. Two platforms are working; six more are designed. No GPU. No floating point. No cloud.

04 What we learned — including the honest edges

It uses the *public* GRU, on purpose. This deliberately runs the well-known GRU (Cho 2014), **not** the lab's proprietary wave architecture. That is a two-tier IP strategy in action: a **public face that can go viral without leaking the private core**. Saying so is part of the honesty.

It is a floor, not a ceiling. This is a proof of *possibility* and a demo — not a useful model. It shows what is minimally required, which is the interesting number, not what is maximally capable.

"Fakes don't run on a Z80." The reason this is such a good demo is that the hardware is the referee. There is nowhere to hide a fudged result inside 2.2 KB and a 3.5 MHz clock.

Why give this away?

Some results are worth more as gifts than as secrets. A 2.2 KB AI on a 44-year-old machine is a front-page story that costs us nothing to share, because it is built on public science and cannot be monetised directly. Open the wonder; keep the machine. This is the wonder.

05 Where it went / status

Kept **open**, as a demonstration and a recruiting tool for the lab's core belief: intelligence you own should not require hardware you rent. It pairs with the sub-£600 BCI at the sensing end — cheap sensors in, tiny models on the metal — to make one argument, from two directions, that the floor is far lower than the industry assumes.

What is still open — kept visible

The honest edges of this lane, shown next to the wins. None of these is a footnote; each is a claim we have not yet earned — and marking it is what earns the right to be believed on the rest.

- ☐ It uses the public GRU (Cho 2014), NOT our proprietary architecture — a deliberate public face that can go viral without leaking the core IP.
- ☐ Two vintage platforms are working; six more are designed, not done.
- ☐ It is a proof-of-possibility and a demo, not a useful model — the point is the floor, not the ceiling.

06 Sources, glossary & the ledger

Where every claim comes from — and what the plain words mean.

Sources

- 01 **Vocabotics project audit — Retro / 'AI on every computer ever made' (2.2 KB GRU on Z80), Apr 2026** — Vocabotics internal project history (as of April 2026)
- 02 **Cho et al., Gated Recurrent Unit (the public architecture used for the demo)** — the public model this demo deliberately uses (as of 2014)

THE HONESTY LEDGER

- Shipped — measured, gate-clean.
- 🔬 Research — measured, partial; the gap is shown.
- 🔭 Vision — designed, honest; not yet built.

The bigger claim goes up a tier, never down. We show the eval beside the demo — the tiering is the credibility.

See the open BCI at the sensing end

Cheap sensors in, tiny models on the metal — the same argument that owned intelligence does not require a datacentre.

[/lab/sub-600-bci](#)

Nothing faked.

A real, dated lab report from our own R&D. Measured, not asserted; the honest negatives are kept on the page next to the wins — because a lab that hides its failures forfeits the right to be believed on its results.

