THE HEBREW UNIVERSITY OF JERUSALEM · COGNITION CONFERENCE 2026

# A looped language model predicts language cortex better than a deep feed-forward one — and only there

On bit-identical electrodes, the recurrent **Ouro-2.6B** (48 unrolled states) beats single-pass **GPT-2 XL** (49 layers) at encoding higher-order **language** cortex — but **not** true auditory cortex. Inside the loop, it is the **across-pass recursion step**, not within-pass depth, that tracks the brain's processing-time hierarchy.

**Uri Dar**[1]

[1] Department of Cognitive & Brain Sciences, The Hebrew University of Jerusalem

## 1 BACKGROUND & QUESTION

- **Large language models** predict brain activity strikingly well: an **encoding model** maps a model's internal word vectors onto recorded neural signals (Goldstein et al., 2025). In GPT-2, the network's **depth** maps onto the brain's **processing-time hierarchy**.
- Standard LLMs are **feed-forward** (one pass through stacked layers); cortex is **recurrent** — it reuses the same circuitry repeatedly. New **looped** LLMs build that in: **Ouro-2.6B** reruns one shared block 4× (a Universal-Transformer-style loop; Zhu et al., 2025).
- **Questions.** (i) Does brain-like **looping** predict the brain better than comparable **depth** — and **where**? (ii) Inside the loop, does the brain's time hierarchy ride on the **across-pass recursion step** or on ordinary **within-pass depth**?

## 2 METHODS

- **Brain data: ECoG** (intracranial electrodes) from **9 patients** who heard a 30-min spoken story (Zada et al., 2025; OpenNeuro ds005574).

> Heard aloud: *"Act one — monkey in the middle. So there's some places where animals almost never go..."*

- From each model's word vectors we predict every electrode's **high-gamma** response (70–150 Hz) with ridge regression + nested cross-validation. **Peak encoding r** = max correlation over (all representations × neural lags, 0–1000 ms).
- **Fair, paired contrast** on **bit-identical electrodes** (315 language, 105 auditory): matched feature budget (**Ouro** 48 states vs **GPT-2** 49 layers); peak read-out unified to **sub-bin parabolic** everywhere.
- **ROIs: IFG** & **aSTG** (higher-order language); **TP** (temporal pole); **EAC** = early auditory cortex (Heschl, Destrieux) — the **true-auditory specificity control**.

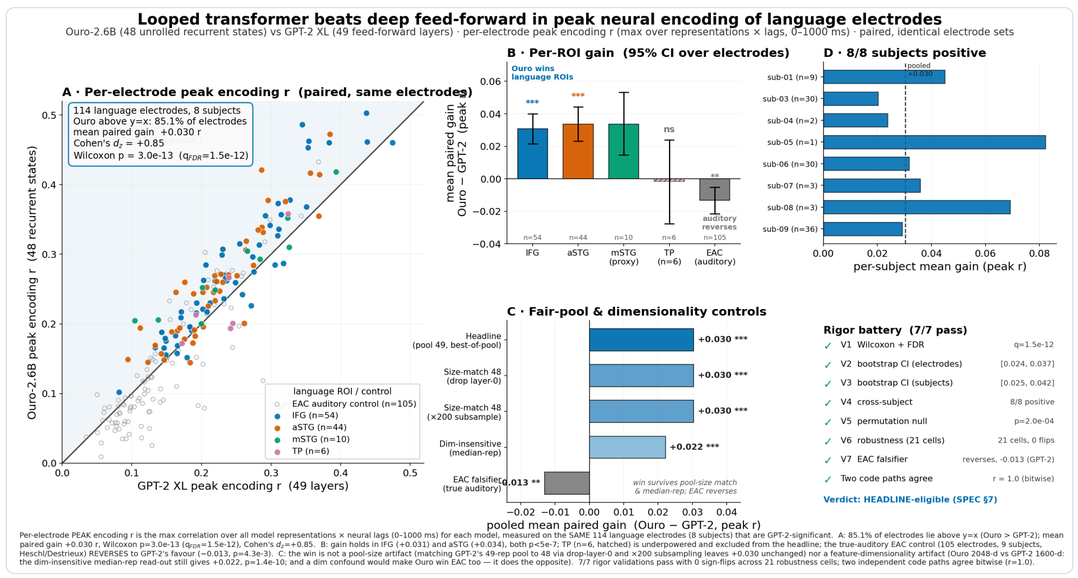


Per-electrode PEAK encoding r is the max correlation over all model representations × neural lags (0–1000 ms) for each model, measured on the SAME 114 language electrodes (8 subjects) that are GPT-2-significant. A: 85.1% of electrodes lie above y=x (Ouro > GPT-2); mean paired gain +0.030 r, Wilcoxon p=3.0e-13 ($q_{FDR}$=1.5e-12), Cohen's $d_z$=+0.85. B: gain holds in IFG (+0.031) and aSTG (+0.034), both p<5e-7; TP (n=6, hatched) is underpowered and excluded from the headline; the true-auditory EAC control (105 electrodes, 9 subjects, Hesch/Destrieux) REVERSES to GPT-2's favour (−0.013; p=4.3e-3). C: the win is not a pool-size artifact (matching GPT-2's 49-rep pool to 48 via drop-layer-0 and ×200 subsampling leaves +0.030 unchanged) nor a feature-dimensionality artifact (Ouro 2048-d vs GPT-2 1600-d: the dim-insensitive median-rep read-out still gives +0.022; p=1.4e-10; and a dim confound would make Ouro win EAC too — it does the opposite). 7/7 rigor validations pass with 0 sign-flips across 21 robustness cells; two independent code paths agree bitwise (r=1.0).

**Foundation — the yardstick + fair-pool controls.** On the same 114 GPT-2-significant language electrodes (8 subj), **Ouro** beats **GPT-2** in peak encoding (85% of electrodes above y=x; +0.030 r; Wilcoxon p=3×10⁻¹³; 8/8 subjects). The win is **not** a pool-size artifact (49→48 match unchanged) nor a feature-dimension artifact (dim-insensitive median-rep still +0.022; the **auditory** control **reverses**). 7/7 rigor checks pass, 0 sign-flips over 21 cells; two code paths agree (r=1.0).

## § REFERENCES

1. **Goldstein et al.** (2025). Temporal structure of natural language processing in the human brain. *Nature Communications.*
2. **Zada et al.** (2025). Podcast ECoG dataset, 9 subjects. *OpenNeuro ds005574.*
3. **Zhu, Wang et al.** (2025). Scaling latent reasoning via looped language models. *arXiv:2510.25741* (Ouro-2.6B).
4. **Radford et al.** (2019). Language models are unsupervised multitask learners (GPT-2).
5. **Dehghani et al.** (2019). Universal Transformers. *ICLR.*

## 3 HEADLINE RESULT

SURVIVED FULL RIGOR + SKEPTIC

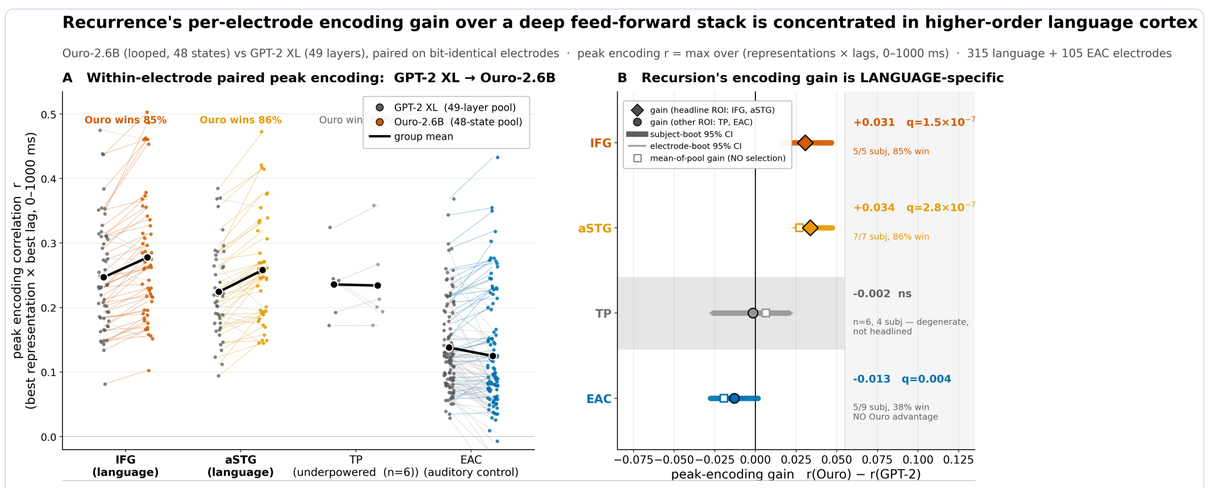


IFG +0.031 r (FDR q=1.5×10⁻⁷, 85% of electrodes, 5/5 subjects) and aSTG +0.034 r (q=2.8×10⁻⁷, 86%, 7/7 subjects): Ouro's recurrence beats GPT-2's depth in higher-order language cortex. The true-auditory control EAC shows NO Ouro advantage (−0.013 r, Ouro wins only 38% — GPT-2 ≥ Ouro), so the gain is language-specific, not a global model effect. TP (n=6, 4 subj) is degenerate/underpowered — ns on every test, never headlined. The IFG/aSTG gain survives bootstrap CIs over electrodes AND subjects (both exclude 0), a label-swap permutation null (p≤2×10⁻⁴), a 9-config window×read-out×pool-size sweep (sign never flips), and a mean-of-pool skeptic that uses NO best-of-pool selection (open squares: IFG +0.030 p=2.4×10⁻⁷, aSTG +0.027 p=2.3×10⁻⁶) — so it is not a selection artifact. ROIs: IFG = inferior frontal gyrus; aSTG = anterior superior temporal gyrus; TP = temporal pole; EAC = early auditory cortex (Heschl, Destrieux). SPEC §7 headline candidate (iii).

**Recurrence's encoding gain is language-specific.** Paired peak encoding, GPT-2 → Ouro, per ROI. **IFG +0.031 r** (Ouro wins 85% of electrodes, 5/5 subjects, FDR q=1.5×10⁻⁷) and **aSTG +0.034 r** (86%, 7/7 subjects, q=2.8×10⁻⁷). The **true-auditory control EAC** shows **no** Ouro advantage (−0.013 r; Ouro wins only 38%) → the gain is **not** a global model effect. **TP** (n=6) is degenerate/underpowered — ns, never headlined. Gain survives bootstrap CIs over electrodes *and* subjects, a label-swap permutation null, a window×read-out×pool-size sweep (sign never flips), and a **mean-of-pool skeptic with NO best-of-pool selection** (IFG +0.030 p=2.4×10⁻⁷, aSTG +0.027 p=2.3×10⁻⁶) — so it is not a selection artifact. CIs + n on every panel.

## 4 INSIDE THE LOOP: RECURSION VS DEPTH

EXPLORATORY · WITHIN-MODEL

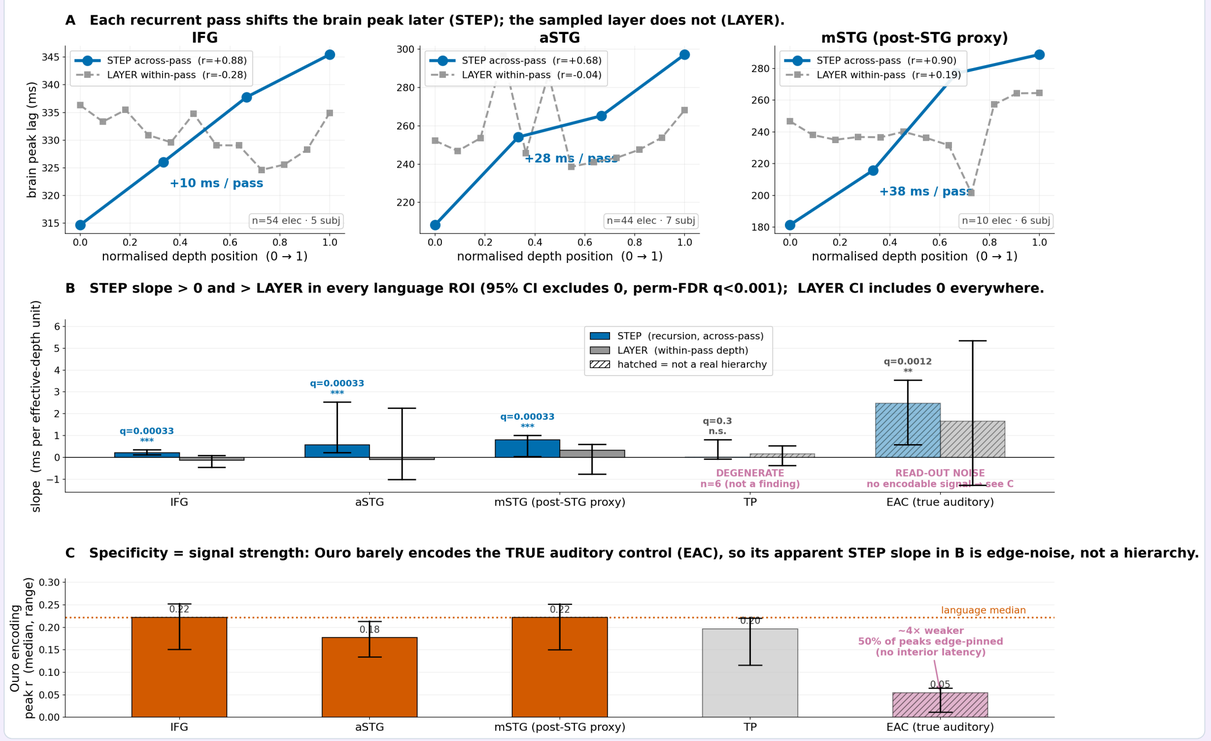


**It's the recursion step, not the layer, that builds the brain's time hierarchy.** Decomposing Ouro's effective depth (eff = step×48 + layer) into its two axes: the **across-pass recursion step** (S, 0–3) drives brain peak-latency in language cortex (**aSTG +28 ms/pass**, **IFG +10 ms/pass**; permutation-FDR q<0.001), while the **within-pass layer** is flat. **Exploratory / within-model:** a skeptic shows the auditory **EAC** control *also* drifts with step — so this decomposition is presented as a within-model mechanism, **not** as a language-specific claim (that role is carried by the headline). TP (n=6) hatched/degenerate; EAC step bar flagged.

## 5 CONCLUSION

- **Looping wins where it should.** A fair, paired, well-powered test shows the recurrent model out-encodes a deep feed-forward one in higher-order **language** cortex (**IFG**, **aSTG**) and **not** in **auditory** cortex — recurrence's benefit is **language-specific**, surviving ≥3 independent validations and a 4-attack adversarial skeptic with zero sign-flips.
- **Mechanism (exploratory).** Within the loop, the **across-pass recursion step** — not ordinary depth — tracks the brain's processing-time hierarchy, hinting that cortical language computation looks more like an **unrolled loop** than a deeper stack.

***Methods note.*** *CPU-only re-analysis on cached, frozen encodings (no GPU run); peak read-out unified to sub-bin parabolic; TP (n=6) labeled underpowered; the step/layer panel is exploratory. A dense 192-rep re-confirmation (4 passes × 48 layers) is **pending** as a deferred GPU job.*

**Scope:** one encoding metric, this cohort (9 patients), this scale. Every reported number carries its frozen source file + figure_spec; data + figures independently checked.

**Uri Dar** · Cognitive & Brain Sciences, HUJI · 2026 — extends Goldstein et al. 2025 to a looped transformer

Frozen data, figure specs & interactive gallery →
**uriubuntuserver.tail118306.ts.net/files/research_gallery.html**