Hippocampal Acetylcholine and Running Speed

Five candidate functions, six data segments, and a one-second temporal lead

Kai Gohil, Laidlaw Scholar, MSci Neuroscience, UCL

Supervisor: Prof. Caswell Barry, UCL Cell & Developmental Biology

Why running speed?

de Cothi, Shipley & Barry (2026, *Nat. Rev. Neurosci.*) propose that hippocampal ACh encodes a **state-transition prediction error (STPE)**, gating plasticity analogously to dopamine's role in reward prediction error. The theory predicts a positive ACh–speed relationship, since unresolved transition error accumulates faster the faster an animal moves.

Dataset: GRAB-ACh fibre photometry (470 nm signal, 405 nm isosbestic control) from a single head-fixed mouse across four VR-track sessions.

Caveat: one animal across four sessions. Generalisability across animals is not established.

Methods & behavioural paradigm

- Raw, unbinned ΔF/F and speed arrays throughout; never binned or smoothed prior to fitting.
- Five candidate functions compared via AIC/BIC computed from first principles (RSS), never a packaged stats function.

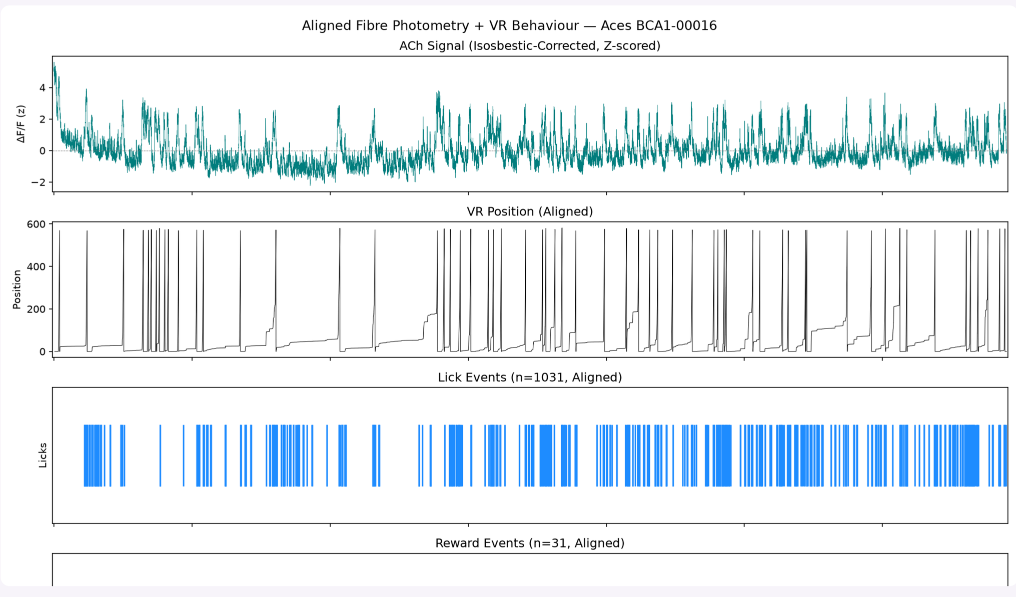

Day 1 session: ACh signal aligned with VR position and lick events (n=1,031) across ~3,400s. Illustrates the recording paradigm.

The baseline speed–ACh relationship

On Day 1 (most extensively analysed; isosbestic R^2=0.978, n=12,154), **log-quadratic** wins decisively by AIC/BIC (R^2=0.453). Fitted form $a \cdot \ln(x)^2 + b \cdot \ln(x) + c$ with $a < 0$; turning point ≈508 pos-units/s, well outside Day 1's observed range (~250 max).

Days 2–3 initially fit poorly, traced to a discrete mid-session fibre-coupling shift. Splitting each session at its bimodal step recovered the relationship (Day 3 post-shift R^2: 0.024 → 0.411).

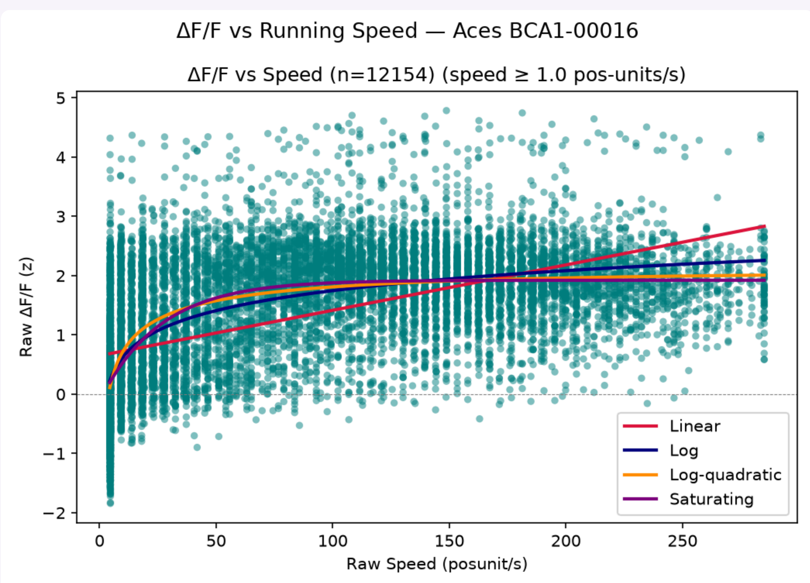

Four of five candidate functions (power-law omitted; its unconstrained fit collapses to a disguised log model) fit to raw Day 1 data (n=12,154). Log-quadratic and saturating exponential track almost identically; linear (red) misses the plateau at both ends.

What this could mean

The lag is the most scientifically interesting result of the placement, for two reasons beyond its statistical strength.

1 It's in direct tension with a naive reading of STPE. If ACh reports an error computed from an already-experienced transition, a same-time or lagging relationship with speed would be more natural. A one-second *lead* instead raises the possibility that the signal reflects something closer to an upcoming motor or behavioural state — not, or not only, a retrospective error report.

2 Two candidate readings remain compatible with the data.

- **ACh encodes a prediction** of the next state, rather than the error on the last one.
- **ACh reflects motor intention/preparation**, distinct from prediction error altogether.

Distinguishing them is the natural next step (see caveat, top-right, and the acceleration analysis opposite).

ACh changes before speed does

Cross-correlating raw, unsmoothed ΔF/F against raw running speed recovers a negative peak lag of **−1.0 to −1.2 seconds** (negative = ACh precedes speed). Consistent across all six data segments (two full sessions plus Days 2 and 3 each split at a mid-session probe-coupling artefact).

Lag-correction improves the best model's fit in four of six segments (Day 1: R^2=0.506 → 0.595, same filter); the two noisiest pre-shift segments show no reliable change.

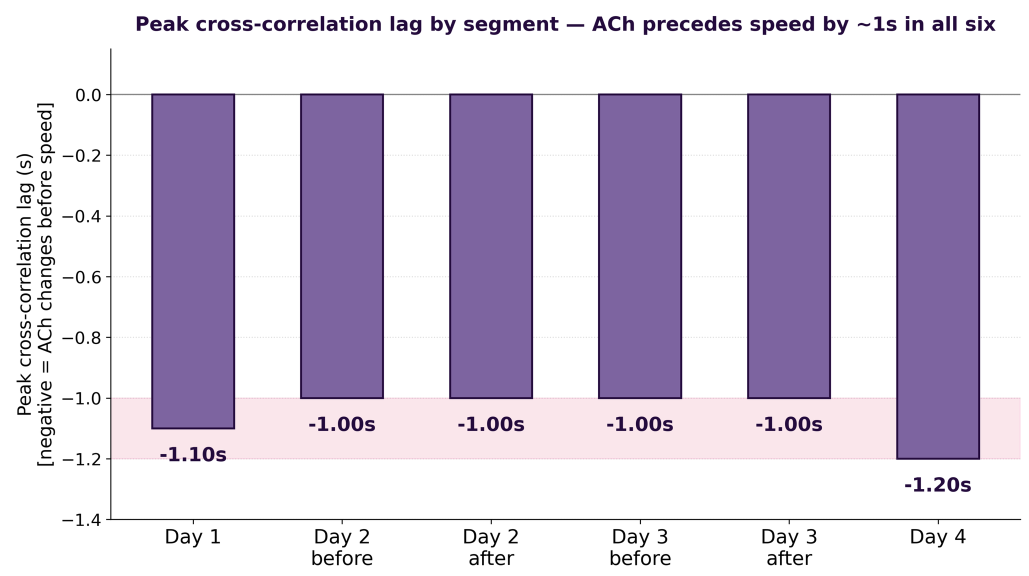

Peak cross-correlation lag values for six data segments; summary of six numbers, not a continuous correlogram trace.

Sync-alignment check (below) rules out clock drift as the source of the lead. A fixed sensor-kinetics contribution remains to be tested via a synthetic sign-flip.

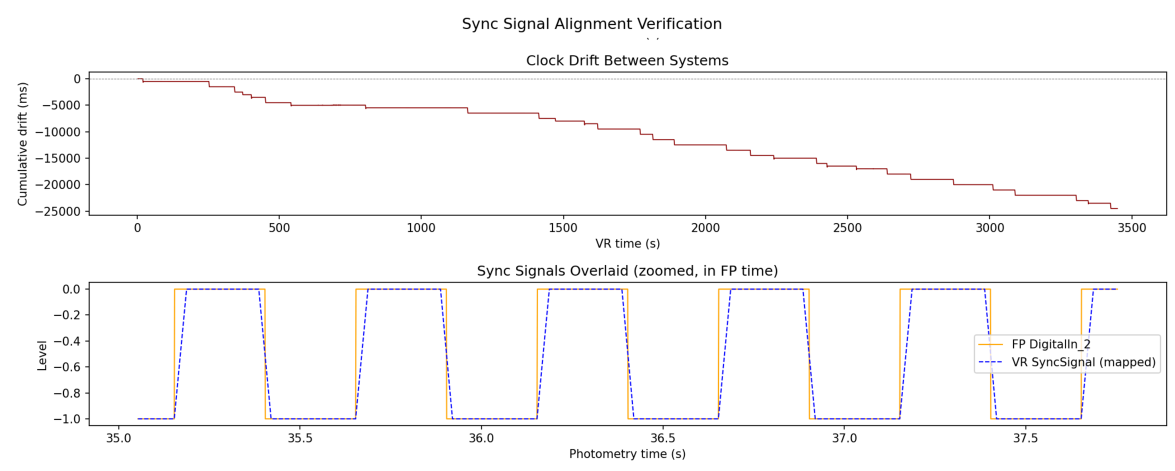

Photometry–VR sync verification, Day 1: raw cumulative clock drift (top) is corrected via a linear fit of matched sync edges. After mapping, sync signals overlay within tens of ms (bottom) — well below the ~1s ACh-precedes-speed lead.

What about acceleration?

Acceleration is a much weaker predictor than speed (R^2≈0.13 vs ≈0.45–0.51) and adds negligible power on top of speed (tested additively across all six segments). It also shows no comparable stable temporal relationship — peak lags ranged +1.5s to −0.1s with no consistent direction. Two speed×acceleration interaction terms (prompted by a mid-placement AI-assisted LLM review) neither improved the fit.

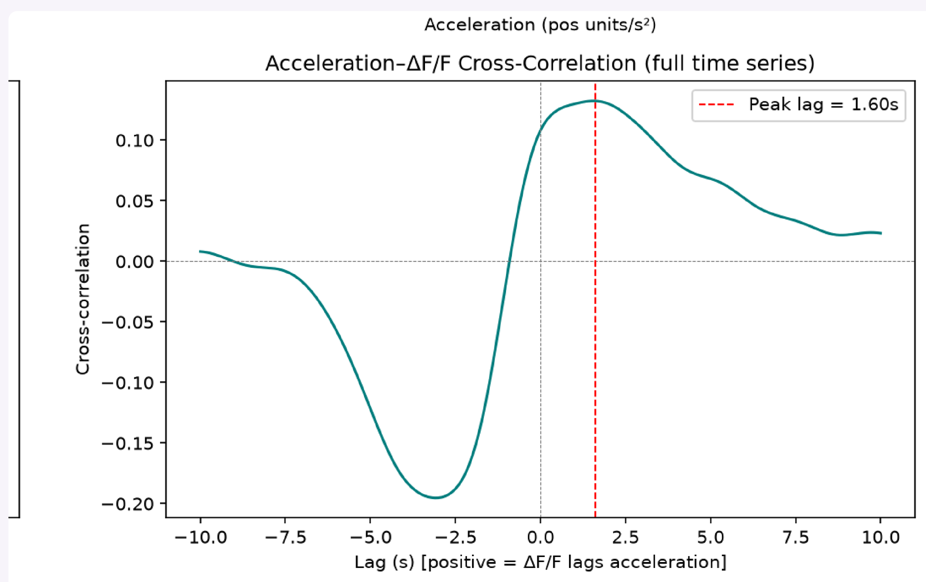

Day 1 acceleration–ΔF/F cross-correlation: peaks at +1.60s — opposite sign and different magnitude from speed's −1.10s. Not stable across segments (range +1.5s to −0.1s).

"Only speed shows a stable timing relationship with ACh."

What's next

- Confirm the lag with a single identical-filter rerun of all six segments.
- Complete the sensor-kinetics ruling-out with a synthetic sign-flip test.
- Check Day 4's log-quadratic turning point (≈61 pos-units/s, coefficients a≈−0.258, b≈2.125, c≈−2.688) against its observed speed range.
- Distinguish the two interpretations opposite by event-locking to specific predictable transitions.

References: de Cothi, W., Shipley, S., & Barry, C. (2026). Acetylcholine: a candidate substrate for hippocampal predictive learning? *Nature Reviews Neuroscience.* doi:10.1038/s41583-026-01058-w.

Acknowledgements: Thanks to Prof. Caswell Barry and Dr Sarah Shipley, UCL. Funded by the Laidlaw Scholarship.

Attribution: Analysis code extended a base pipeline written by Dr Sarah Shipley (basicPlotting.py, fp_utils.py); model-comparison, session-splitting, acceleration and cross-correlation scripts written by Kai.