Resubmission cover letter

11 September 2026

Dear Editors-in-Chief and Special Call Editors,

Please reconsider “TaxBridge-Bench: A Neurosymbolic Benchmark for Spreadsheet-to-Ledger Reasoning in Tax Compliance” as a resubmission of tracking number **1003-2039**, previously recorded as “Article in Special Issue”, for the **Special Call for Neurosymbolic Benchmark Papers**.

Thank you for identifying the missing author name and unfinished data-availability statement. I have addressed both points and reviewed the manuscript against the journal's instructions.

1. **Author identification.** David Elliman now appears on the first page with affiliation and correspondence details. The running head and PDF author metadata are also corrected.
2. **Data availability.** The placeholder has been replaced by the verified public archive DOI: https://doi.org/10.5281/zenodo.22707695. TaxBridge-Bench version 1.0.0, published on 11 September 2026, contains the 30 evaluation workbooks, 40 deployment workbooks with executable specifications and ground truth, 45 diagnostic workbooks, scorer/harness source snapshots, saved baseline prompts and responses, and recovered comparison records. Data and documentation are CC BY 4.0; the eight source-code files are MIT-licensed. The archive provides licence scope, an inventory and checksums. Public access and the downloaded archive's checksum have been verified. The manuscript distinguishes this dataset from the separate architecture preprint and discloses unavailable records and runtime dependencies.
3. **Manuscript review.** The revision follows the Royal single-column format, has a 158-word unstructured abstract and six keywords, and remains 12 pages. References have been checked, the dataset has been cited, and the required declarations added. Funding is stated as Neuro-Symbolic Ltd's own resources, with no external funding received.

Checking the recovered records also led to clarifications of the scoring criteria. The paper distinguishes aggregate-box passes from row-level money, period and VAT release checks, where category differences can remain advisory. Re-scoring the saved simple-baseline responses reproduces the reported box and flag results. Unsupported full-corpus blackboard token estimates have been removed, and diagnostic and ablation claims lacking matching records are identified as unverified. No new model experiments were run for this revision.

My related preprint, “Bounded LLM Reasoning: A Neuro-Symbolic Blackboard Architecture for Production Document Analysis” (version 1.1, https://doi.org/10.5281/zenodo.21426202), describes the reference architecture and earlier evaluations. The present manuscript cites it and focuses on the benchmark task, datasets and scoring, including the deployment corpus and serialised-workbook baseline. This overlap is disclosed for editorial assessment.

I am the founder of Neuro-Symbolic Ltd and developer of the reference implementation. The study uses synthetic data and includes no raw customer or human-participant records. OpenAI Codex assisted with this revision, reference checking, the audit of saved experimental records and preparation of these materials; I remain responsible for their content.

Thank you for considering the revised manuscript.

Sincerely,

David Elliman
Neuro-Symbolic Ltd, Gloucestershire, United Kingdom
dave@neusym.ai