

A Bounded Neuro-Symbolic Loop with EML Trees and qNS Masks

Dana Edwards


April 20, 2026

Abstract

This note records a bounded Tau-adjacent experiment in which a neural or external proposer emits symbolic formula candidates, a small EML grammar gives the candidate language, a finite qNS Boolean algebra routes evidence masks, and Tau checks promotion and table-memory updates over finite audited masks. The central pattern is

$$q_{\text{NS}}(y \mid x) \propto q_{\text{N}}(y \mid x)\chi_{\text{S}}(y, x),$$

read as: the neural side proposes and scores, while the symbolic side filters. The current implementation is not native analytic Tau semantics and not a general symbolic-regression theorem. It is a certificate-carrying bounded loop: formula proposal, parse and domain checks, finite sample checks, proof metadata, qNS mask routing, tamper rejection, and pointwise memory revision.



1. qNS as a finite Boolean algebra. For an audited candidate universe U , the qNS carrier is

$$B_U := \mathcal{P}(U).$$

Thus $0 = \emptyset$, $1 = U$, $A \wedge B = A \cap B$, $A \vee B = A \cup B$, and $A' = U \setminus A$. In the Tau experiment this is represented as an eight-bit carrier `qns8`. A hard filter has the exact form

$$q_{\text{NS}}(c \mid x) = \frac{q_{\text{N}}(c \mid x)\chi_{\text{S}}(c, x)}{\sum_{d \in C_x} q_{\text{N}}(d \mid x)\chi_{\text{S}}(d, x)}.$$

The formula does not say that Tau computes neural probabilities. It says that after the host proposes and scores candidates, the symbolic layer deletes candidates whose indicator is false and the host renormalizes the remaining mass.

2. EML as the symbolic hypothesis language. Following Odrzywolek's elementary-function basis proposal, the experiment uses the grammar

$$T ::= x \mid 1 \mid \text{eml}(T, T), \quad \text{eml}(a, b) = \exp(a) - \ln(b).$$

For data D , the candidate-tree filter is

$$q_{\text{NS}}(T \mid D) \propto q_{\text{N}}(T \mid D)\chi_{\text{parse}}(T)\chi_{\text{domain}}(T, D)\chi_{\text{spec}}(T, D)\chi_{\text{proof}}(T).$$


The specification gate used by the bounded regression demo is

$$\chi_{\text{spec}}(T, D) = 1 \iff \text{MSE}_D(T, f) \leq \varepsilon.$$

Standard reading: the candidate survives the specification gate exactly when its mean-squared error against the target on the declared finite data set is at most the tolerance. This is a bounded data check, not a global identity theorem.

3. Certificate-carrying promotion. The implemented promotion rule is the conjunction of explicit evidence gates:

$$\text{Promote}(T) \iff \chi_{\text{parse}}(T) \wedge \chi_{\text{domain}}(T, D) \wedge \chi_{\text{spec}}(T, D) \wedge \chi_{\text{proof}}(T) \wedge \chi_{\text{obligation}}(T).$$



Each gate has a different purpose. Parsing prevents the proposer from escaping the grammar. Domain checks prevent invalid real logarithms on the bounded interval. Specification checks compare finite samples. Proof metadata records whether the current survivor shape is inside the known certificate surface. Obligation checks connect constructor-plan side conditions to sidecar interval evidence.

The current bounded regression surface enumerates $\mathcal{T}_{\leq 3}$, a depth-three EML corpus of 1446 trees, and selects

$$T^* = \min_{\prec} \{T \in \mathcal{T}_{\leq d} : \text{err}_{\text{train}}(T, D) \leq \varepsilon \wedge \text{err}_{\text{holdout}}(T, D) \leq \varepsilon\}.$$

In the checked demo corpus, the selected formulas include the compact representatives for x , $\exp(x)$, $\ln(x)$, and $\exp(\exp(x))$. The noisy regression artifact records signed residuals rather than claiming a general recovery theorem.

4. Tau-facing memory update. Accepted sidecar evidence is converted into qNS masks and routed through Tau. The table-memory state is a finite map

$$M : I \rightarrow B_{\text{qNS}}.$$

The pointwise update is

$$M_{t+1}(i) = (G(i) \wedge A(i)) \vee (G(i)' \wedge M_t(i)).$$

Standard reading: for each memory slot i , the next memory value is the join of the replacement value restricted by guard $G(i)$ and the old memory value restricted by the prime of $G(i)$. In the Tau experiment this appears as a named table-backed revision:

```
memory_revise_qns8(old, guard, replacement) ≡ table {when guard => replacement; else => old}.
```

5. Current executable receipt. The public demo path in `TauLang-Experiments` downloads official Tau, applies local research patches, builds Tau, and runs the qNS/EML demos. The current model-free memory fixture reports:

```
candidate count = 5,
parse-ok count = 3,
promoted count = 3,
review count = 2,
memory-updated count = 3,
rejected-preserved count = 2.
```

It also checks that qNS table regressions and symbolic Tau table checks pass. The important result is not that a language model is trusted. The result is that model or fixture proposals can enter a narrow evidence pipeline, and only checked masks can update memory.

Boundary. The experiment does not prove full symbolic regression, complete EML normalization, native Tau evaluation of \exp or \ln , complex principal-branch correctness, or production safety of an attached LLM. The mathematical contribution is a scoped architecture:

Proposal \rightarrow Grammar \rightarrow FiniteChecks \rightarrow ProofMetadata \rightarrow qNSMasks \rightarrow TauGate \rightarrow MemoryRevision .

This is useful because every promotion claim is reducible to finite artifacts that can be replayed, inspected, and rejected if stale or tampered.

References. Andrzej Odrzywolek, “All elementary functions from a single operator,” arXiv:2603.21852v2, 2026. Ohad Asor, *Theories and Applications of Boolean Algebras*, draft v0.25, 2024. TheDarkLightX, *TauLang-Experiments*, public experiment repository, [GitHub](#). Dana Edwards, “Neuro-symbolic Boolean algebras in Tau Language,” Tutorial 49, 2026. Dana Edwards, “EML trees as neuro-symbolic hypotheses,” Tutorial 50, 2026. Dana Edwards, “Symbolic hypothesis generation with EML and qNS,” Tutorial 51, 2026.