Delta

An AI Guided Knowledge Network for Explanation Verified Understanding

"What I cannot create, I do not understand."

Richard Feynman [6]

Abstract. Delta is an AI guided learning system that turns a learner's own study material into a structured knowledge network, and only extends that network once the learner has explained a connection well enough for an AI to verify it. The design rests on three principles: the network is grounded in material the learner already trusts, no connection counts until it has been actively derived or explained, and a failed explanation routes back to the nearest already confirmed piece of knowledge rather than forward to a new, isolated exercise. The result is a map that reflects what a learner actually understands, not what they have merely read.

---

1. The Problem

Knowledge is the possession of information, and a purely additive collection of information can be sufficient when that information has no inner structure, such as vocabulary, names, or dates, because nothing needs to be derived from it.[1] When information instead carries an underlying causal or logical structure, a memorized collection is unsatisfying in the long run: understanding becomes possible only once several elements are connected in a way that lets a learner explain and derive the relationships between them.[1][2] Material learned by rote stays isolated in hierarchically organized folders, fades once it is no longer used, and cannot be reconstructed once it is forgotten.[5]

Conventional assessment reinforces the shallow path. Examination formats reward short term recall over derivable understanding, a pattern documented as the backwash effect: learners study toward whatever the test actually measures, not toward the structure the material contains.[4] Standard diagnostic instruments in physics show the same gap directly, students who can recite a formula frequently cannot apply the underlying concept once the surface wording of a problem changes.[3] No common tool currently tests, at the moment of learning, whether a connection has been understood rather than only encountered.

2. Core Design Principles

Material grounded network. The knowledge network is decomposed by an AI directly from material the learner submits, such as lecture notes, textbook sections, or problem sets, rather than generated from a pre built subject corpus. Structure always reflects what the learner is actually studying, not a generic curriculum.

Explanation gated connections. A connection between two knowledge elements, for instance a derivation, a special case, or an application, is proposed by the AI but only marked understood once the learner supplies an explanation in their own words or their own calculation steps, and the AI verifies it. Passively accepting a suggested link changes nothing in the network's state.

Backtracking to secured knowledge. When an explanation fails verification, the system does not present an unrelated new exercise. It moves the learner to the nearest node that is already confirmed and restarts derivation from there, rebuilding the chain step by step until the original connection can be reached honestly.

3. Technical Architecture

Submitted material is first decomposed by an AI into discrete knowledge elements, the nodes of the network. The AI then proposes candidate edges between nodes, derivations, special cases, and applications, each carrying a status of collected, proposed but not yet explained, or confirmed, actively explained and verified. Every edge the learner works through moves through a verification step before its status can change.

The network is rendered as a navigable three dimensional space rather than a flat list or a strict hierarchy, extending the underlying idea of a concept map into an explicit, explorable layout.[2] Distance from the center encodes closeness to foundational principles, relative position encodes thematic relatedness between elements, and color distinguishes confirmed understanding from knowledge that has only been collected. A learner can see, at a glance, which parts of a subject are structurally secured and which are still loose ends.

4. The Verification Mechanism

Verification follows the same test Feynman applied to his own understanding, that a concept is only known once it can be reconstructed and explained from first principles in one's own terms.[6][7] The learner is asked to reconstruct the reasoning behind a proposed connection, in prose or in calculation, rather than to select or confirm an answer the AI already produced. In STEM material this reconstruction can additionally be checked against the formal derivation, giving the AI an objective pass or fail criterion. That objective check is also why STEM material is the starting domain: it gives verification a ground truth before the same explanation gated approach is extended to fields where an explanation can only be judged for logical coherence rather than formal correctness.

5. Why Structure Produces Understanding

Two established findings motivate the design. First, retrieval, actively reconstructing information from memory, produces more durable learning than re reading or passive review, an effect long established in memory research and known as the testing effect.[8] Requiring an explanation before a connection counts as understood embeds this retrieval step directly into how the network is built, rather than leaving it to separate, optional revision. Second, connected, elaborated knowledge resists forgetting better than isolated facts, because a forgotten element can be re derived from neighboring elements instead of being permanently lost.[2][5] A spatial network in which every edge has to be earned through explanation is, in effect, a continuously updated record of exactly which parts of that structure a learner can currently reconstruct.

6. Limitations & Open Questions

AI verification of an explanation can misjudge in both directions, accepting a superficially correct but shallow explanation, or rejecting a correct explanation phrased unconventionally; the objective check available in STEM material narrows this risk but does not remove it in less formal domains. The system also depends on material that has enough internal structure to be worth connecting at all, purely additive information such as vocabulary gains little from an explanation gate.[1] Finally, every network starts empty: a learner must submit and decompose material before any structure exists, so the approach is best suited to learners already engaged with a body of material rather than to a cold start with no prior content.

7. Long-Term Vision

The aim is a map of understanding a learner can actually trust, one where every visible connection has been earned by explanation rather than assumed from exposure. Starting in STEM material, where derivations can be checked formally, establishes that trust before the same explanation gated, spatially structured approach extends to fields where understanding can only be judged by the coherence of the reasoning itself. What cannot yet be created from first principles is, by that standard, not yet understood, and Delta is built to make that standard visible while it is still being met.

References

[1] Kvanvig, Zagzebski, Carter, and Gordon, "Understanding," Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/entries/understanding/

[2] Ausubel, summarized in "Theory of Concept Maps," IHMC. https://cmap.ihmc.us/docs/theory-of-concept-maps

[3] Hestenes, Halloun, and Swackhamer, Force Concept Inventory. https://en.wikipedia.org/wiki/Force_Concept_Inventory

[4] Biggs, the backwash effect, summarized by Queen Mary Academy.
https://www.qmul.ac.uk/queenmaryacademy/educators/resources/curriculum-design/constructive-alignment/

[5] Willingham, "Do Students Remember What They Learn in School," AFT. https://www.aft.org/ae/fall2015/willingham

[6] Feynman's last blackboard at his Caltech office, documented by Fermat's Library. https://x.com/fermatslibrary/status/1821965763365732788

[7] Feynman, chapter on his time in Brazil, "Surely You're Joking, Mr. Feynman!" https://dhruvdh.github.io/ni/appendix.A.html

[8] Testing effect, retrieval practice and long term retention. https://en.wikipedia.org/wiki/Testing_effect