

The Maze Transformer: Mapping the Claims of a Research Field onto the Vacuum.Net

An instrument that reads scientific papers, keeps only their quantitative claims, chains those claims across papers and tests the chains. This essay states its purpose, method, contents and results, and how it is linked to both MAZEs.

J. Konstapel, Constable Research, Leiden – with computational assistance from Claude (Anthropic) · 28 October 2026 · live: maze.swarp.nl/transformer · code and plans: [maze/transformer](#) in the SWARP repository

Abstract. A research field states its knowledge in many papers. Each paper makes a few hard, quantitative claims. Few papers state how their claims connect to those of other papers. The Maze Transformer extracts such claims and places them in one space. In the physical sciences a claim is a power law, $y \propto x^e$. In medicine it is an effect size, such as an odds ratio with its interval. Two claims that share a quantity can be chained. A chain predicts a relation that no single paper states. When a third paper measures that relation directly, the prediction is tested.

The instrument rests on the Vacuum.Net theory. In that theory a law between two magnitudes is a straight line in logarithmic position space, and a direction is a trit: +1, 0 or −1. Three maps are live: galaxies (150 arXiv papers), the life sciences (495 arXiv papers) and Parkinson's disease (1,303 PubMed records). In medicine 36 of 41 testable chains were confirmed by an independent measurement ($p = 0.0015$; 24 of 26 under a stricter independence rule, $p = 0.002$). In the two physics maps no paper measures a chained pair directly, so their 4,860 chains remain untested predictions. Where a sure match exists, each quantity is now linked to the existing knowledge net (the old MAZE) and to the derived library (the new MAZE); the papers themselves have become residents of the net. A blind check found 36 of 40 links correct and none wrong.

1. Purpose

The Maze Transformer has one aim. It should show what a research field claims, where those claims meet, and where a claim is still missing.

A field publishes thousands of papers a year. A reader sees them one at a time. Connections between papers live in the heads of experts and in review articles. Those reviews are slow and selective. The Transformer makes the connections explicit and countable.

It does four things:

1. **Extract.** It reads each paper and keeps only its quantitative claims: laws and effects.
2. **Place.** It gives every quantity a place. Papers that use the same quantity meet at that place.
3. **Chain.** It combines two or three claims from different papers into a path. A path predicts a relation that none of the papers states.
4. **Test.** When a further paper measures the predicted relation directly, the path is confirmed, undecided or contradicted.

The result is a visual map of a field. Dots are quantities. Lines are claims. The list of untested paths is a list of concrete research questions. Each question names the two quantities, the predicted exponent or direction, and the papers it rests on.

The design rule is simple. Count what can be counted. Say "untested" when something is untested. Every number on the map can be traced to a sentence or a formula in a named paper.

2. The setting: SWARP and the two MAZEs

2.1 SWARP

SWARP is a family of web applications built by Constable Research in Leiden. Its stated goal is bottom-up: citizens should know what happens in their own neighbourhood and find each other there. Around that goal sit several standalone services. Among them are a daily guide, a local-environment monitor, a financial weather instrument and a set of research tools. The research tools share one knowledge net, called the MAZE, at `maze.swarp.nl`. The Transformer is one of these tools. It runs on the same server as the MAZE and uses its database.

2.2 The old MAZE: a self-addressing net

The MAZE is a knowledge net in which every item has an address computed from its content. It holds about 8.2 million items, called residents. The largest groups are 3.6 million English Wikipedia articles and 4.6 million European patent applications. Smaller groups include mathematical theorems, works of art, music, blog essays and open questions.

An address is a word of 24 trits. Two generations of addressing exist side by side. The first computes the address from the words of a description. The second computes it from meaning: a text embedding is projected on fixed, variance-ordered axes and a ternary tree that adapts to the residents below each prefix. Items with similar meaning therefore share a long address prefix. Five invariants govern the net. The most important are that an address follows from content alone and that an address, once written, is never rewritten.

2.3 The new MAZE: a derived library

The new MAZE is built the other way round. It does not copy any existing classification. It unwinds from nothing, one operation at a time. Each node follows from earlier nodes by one stated operation and carries an exact check,

computed with integer arithmetic or SymPy. If one check fails, nothing is written.

The library now holds 635 nodes in 71 windings. It runs from the first distinction through quaternions, the Dirac algebra and the fermion spectrum, to the hydrogen atom, the periodic table, chemical binding, the genetic code and further layers. Every node has one of five statuses: adopted, derived, corresponding, tested or open. Its names come afterwards, from the literal first sentence of a Wikipedia article with its revision number. On 19 October 2026 the old MAZE was hung under the new one as a branch, so that every resident can be reached from the first step.

3. The Vacuum.Net theory in brief

The Transformer takes its concepts from the Vacuum.Net theory. The canonical statement is the Foundational Paper, eighth edition, version 8.1 (17 October 2026). The essay *Curves in the Vacuum.Net* (24 October 2026) derives what a path in the net is. This section summarises both, as far as the Transformer needs them. It reports the theory's own claims. Their status is the status the theory itself assigns.

3.1 Five axioms

The theory starts from one object: a strand, an ordered course through a set of places. A loop is a return of the strand to a place it has visited. A place visited twice or more is a knot. Five axioms follow.

- **N1, strand.** There is exactly one strand. Every distinction is a loop.
- **N2, closure.** A loop persists only if it lies in a closed window. What does not close passes on as tension.
- **N3, memory.** At every knot the passages are ordered. The first passage fixes the direction; later passages are read relative to it.
- **N4, context.** Every constant is a function of the strand.
- **N5, scale.** Merged into blocks, the coarse strand obeys N1 to N4 again. What does not close at one scale is material for the next.

3.2 Trits, closure and depth

A direction must differ from its reverse. The smallest cyclic group in which that holds is $\mathbb{Z}_3$. So a passage through a knot has one of three values relative to the first passage: +1 (it leaves the same way), -1 (it leaves the way the first came in) or 0 (another way). These values are trits.

Three trits close when their sum is 0 modulo 3. Of the 27 possible windows, 9 close. Every closed word breaks down into three elementary closures: a standing (0), a meeting ($x, -x$) and a ring (x, x, x). Every integer has exactly one balanced-ternary word; its length is its depth in scales. Meetings double and rings triple. Because $2^a = 3^b$ has no solution except $a = b = 0$, the net never closes completely. It comes closest at depths 1, 2, 5, 12, 41, 53, 306 and further. These are the convergents of $\log_2 3$, the source of the Pythagorean comma.

3.3 What the Foundational Paper adds

Edition 8.1 writes the theory as one chain. It starts from the ground of distinction (Spencer-Brown, Kauffman). It continues through the ternary register: unique addresses, the window register and the two ladders. It then states a knot, a double cover and a conserved Hopf charge. It gives a lever for a free path, $\gamma = -(1 + P)/P$. It writes Dirac's equation in Rowlands' nilpotent form; there Dirac is input, not output. It reads Maxwell's quaternions as two readings of one closure, as a correspondence. Units follow Larson's dimensions. The paper separates proved theorems (T1-T23), correspondences, identifications, five predictions with prior commitments and seven open points. It also keeps a register of every check.

3.4 The part the Transformer uses

The Transformer uses one bridge from this theory to ordinary science. A measured magnitude q is placed by its position $p = \log_3 q$, its depth in scales.

THEOREM 8 (CURVES, TWO LAWS)

A power law $q_j = C \cdot q_i^e$ is a straight line in the plane of p_i and p_j , with slope e . Two laws between the same two magnitudes are identical, parallel, or meet in exactly one point.

A law is therefore a slope. The sign of the slope is a trit. A chain of laws through shared magnitudes is a path in the sense of the theory. Its exponent is the product of the exponents along the way. A path either closes against a direct measurement or carries tension. In the new MAZE this reading is fixed in winding 69. Node `lus_helling` holds the slope trit, with Kepler's third law over eight planets as its calibration (slope 1.4978 ± 0.0007 against the law's $3/2$). Node `lus_overal` holds the rule for which paths close everywhere.

4. How the Transformer works

The Transformer has two layers. Layer A reads formulas. Layer B reads effect sizes. Both build the same structure: places, passages, paths and a gap list.

4.1 Places, passages and paths

- A **place** is a quantity, such as stellar mass, luminosity or the risk of dementia.
- A **passage** is one claim in one paper, from one place to another. In Layer A it carries an exponent e and, when given, an uncertainty u . In Layer B it carries a ratio with a confidence interval and a direction.
- A **path** is a chain of two or three passages from different papers through shared places. Its exponent is the product of the exponents. In Layer B its direction is the product of the directions.
- A path is **tested** when some paper measures its two end places directly. It is confirmed when that measurement agrees within the combined uncertainty. It is contradicted when its sign differs, or when it lies more than three combined uncertainties away. Anything between is undecided. A path without a direct measurement is a **gap**: a prediction and a research question.

4.2 Layer A: formulas from arXiv

For each arXiv paper the Transformer fetches the LaTeX source. Fixed, deterministic rules read the formulas. They recognise forms such as $y \propto x^e$, $y \sim x^e$, $y = C x^e$, products of several factors, a power on the left-hand side, logarithmic forms, exponentials and cycles. They read the exponent, an uncertainty in $\pm$ or parenthesis notation, and the definition of each symbol from the surrounding text.

The rules were extended step by step as blind checks exposed errors. Nineteen reading rules now exist (L1–L19). Examples: a value with units is not a law; nothing is read from figures and tables; "for $t \sim \tau$ " is a condition, not a law. Ten path rules (K1–K10) fix how laws chain. A law may be taken in reverse, with exponent $1/e$. A path has two or three links from different papers. A path may not loop back on itself.

A language model (GPT-4o, temperature 0, cached) does two things only. First, it names a symbol when the text does not define it (rule L10). Second, in the life sciences, it filters out "laws" that are really definitions, notations or parameter values (rule L20). Both steps were tested blind (§6).

4.3 Layer B: effect sizes from PubMed

For medicine the claims are effects, not formulas. The input is a PubMed record in MEDLINE format: identifier, year, MeSH terms and abstract. A deterministic reader finds every odds ratio, hazard ratio or relative risk with its interval in the abstract. Places are the MeSH descriptors of the record. A language model then reads, per effect, which exposure and which outcome are meant, and in which direction. It chooses from the MeSH list of that record. Effects on the same pair are pooled with a random-effects model (DerSimonian–Laird). Each edge gets a pooled estimate, an interval, the identifiers of its papers and a direction trit. A path $A \rightarrow M \rightarrow B$ predicts the direction of $A \rightarrow B$ as the product of its two directions.

4.4 "Have it examined"

Any visitor can submit an arXiv identifier. The server fetches the paper, reads its laws with the same rules and lays it against the map. It reports which quantities the paper shares with the map, which new paths it opens, and whether it measures a gap directly. A paper that measures a gap turns a prediction into a test. Limits keep the cost bounded: at most 40 name questions per paper, 300 per day in total, five papers per visitor per day. A paper already examined costs nothing.

5. What is in it, and what can come in

5.1 Contents today

Map	Source	Contents
Galaxies	150 arXiv papers, astro-ph.GA, February 2025, each with at least one \propto	190 places · 210 laws · 2,648 paths
Life sciences	495 arXiv papers, q-bio and physics.bio-ph, March–November 2025	924 laws read · 202 after filter · 2,212 paths
Medicine	1,303 PubMed records on Parkinson's disease, retrieved 25 October 2026	2,184 statements · 1,791 edges · 51,057 paths

The medicine map shows the strongest part of the net: 224 places, 314 edges and the 41 testable paths. All three maps are browsable at maze.swarp.nl/transformer, [/transformer/levenswetenschappen](http://maze.swarp.nl/transformer/levenswetenschappen) and [/transformer/geneeskunde](http://maze.swarp.nl/transformer/geneeskunde).

5.2 What can come in

- **Any arXiv paper with LaTeX.** The reading rules are field-independent. The naming step currently describes the paper as astrophysics, so names in other fields are less reliable until that is generalised.
- **Any disease or exposure in PubMed.** Layer B needs only a PubMed query. Parkinson's disease was the first domain.
- **Publisher XML.** The specification allows JATS XML with MathML for papers outside arXiv. This route is not yet built.
- **Larger corpora.** A single subfield with about 1,000 papers is the next step for physics. More overlap means more chance that one paper measures what two others imply.
- **Measured values.** Catalogues of measured data, such as stellar or galaxy catalogues, could test physics paths directly instead of waiting for a paper.
- **Directed edges.** In medicine, edges from randomised trials, cohorts and Mendelian randomisation can carry direction of cause. The publication type is already stored.

6. Results

Every measurement below was run under a plan whose hash was committed before counting. Failures are reported, not repaired after the fact. Blind checks were done by a separate assessor, a language-model agent, that saw only the items. Planted errors served as a control that the assessor actually checks.

6.1 Does the space see topics?

The first pilot (25 October) still used words as places, as in the first specification, which was later replaced by quantities. It used 30 papers on two unrelated topics. Papers on the same topic shared 14.0 knots on average, against 5.2 across topics ($p = 0.0001$). The share of +1 trits was 0.089 within a topic against 0.0023 across, a factor of 38. The first run was invalid because of an error in my own control; the second, with a corrected control, passed. A caution remains. A +1 at the word level mostly marks a recurring two-word term. This pilot shows that the space separates topics. It does not yet show that trits add anything beyond word pairs.

6.2 Reading laws and naming quantities

- **Naming was the bottleneck.** With the original definition patterns ("where X is..."), 97% of symbols in 50 papers stayed unnamed. Authors rarely write definitions that way. Two patterns and a language-model step were added.
- **Names, blind.** On the galaxy map 45 of 50 names were judged correct: 39 of 44 from the language model, 6 of 6 from fixed rules. Four of five planted errors were caught.
- **Laws in the life sciences.** The first blind check found only 25 of 60 "laws" real (42%). A third were definitions or notation, such as $\text{curvature} = 1/\text{radius}$ or $X \sim \text{Exp}(\lambda)$. After the language-model filter, a fresh blind sample scored 52 of 60 (87%), with all planted errors caught.

6.3 Paths in physics: untested

The galaxy map has 2,648 paths and the life-sciences map 2,212. In neither corpus does any paper measure the two end quantities of a path directly. The paths are therefore predictions only. This is a property of the corpus, not of the method: papers report laws, not paired measurements of quantities that other papers link. A larger subfield or a catalogue of measured values is needed.

6.4 Effects in medicine: tested

- **Binding exposure and outcome.** Fixed rules bound the right exposure and outcome in only about 25% of effects. An independent second implementation of the same specification agreed with the first on 77% of sentences, but both were mostly wrong on the disputed ones. With the language model reading each effect, the second prompt version scored 54 of 60 (90%) in a fresh blind sample.
- **Paths.** Of 41 testable paths, 36 were confirmed by a direct measurement (88%). A permutation test gave $p = 0.0015$.
- **Stricter count, after the fact.** In 13 of the 41 paths the direct measurement came from the same papers as the chain. One paper confirmed itself. Counting only paths whose direct edge shares no paper with its links gave 24 of 26 (92%), $p = 0.002$. This stricter rule was added afterwards; from now on independence is required in every plan beforehand.

What this shows: when two independent papers report $A \rightarrow M$ and $M \rightarrow B$, the product of their directions predicts the directly measured direction of $A \rightarrow B$ in about nine of ten cases. What it does not show: cause. Many paths mix

populations. "Male sex → Parkinson's → mortality" links the risk of getting the disease to the risk of dying with it. A correct direction is not a mechanism.

7. Relations established with the MAZE

On 28 October 2026 the Transformer was linked to both MAZEs. Before that it ran on the same server but stood apart in content. Two plans governed the work. Plan 1 (hash 89aaba04...) failed its blind check. Plan 2 (hash 92b3f7b0...) passed.

7.1 Quantities to the old MAZE

Each quantity on the three maps is matched to an English Wikipedia article. The article must exist under exactly the quantity's name, with only the first letter capitalised. Or Wikipedia must redirect to it while every content word of the name is kept. A disambiguation page never counts. The matched title is then looked up as a resident of the old MAZE. No statistics and no language model take part.

- **Plan 1 failed.** Without these safeguards, 31 of 40 sampled links were correct (77.5%, threshold 85%). The errors had clear causes: disambiguation pages ("survival", "higher"), titles of works ("football players" led to a painting) and redirects to broader concepts ("time interval" to Time, "body mass" to Human body weight).
- **Plan 2 passed.** In a fresh blind sample, 36 of 40 links were correct (90%). Four were judged too broad and none wrong; all three planted errors were caught. The four broad cases share one limit of a name-only match: the article fits the name but not the field's sense, as with "noise" or "effective temperature" in biophysics.
- **Coverage.** Galaxies 42 of 208 quantities (20%), life sciences 114 of 765 (15%), medicine 106 of 224 (47%). Most unmatched names are compound quantities, such as "dark matter halo mass function", for which no single article exists. Each unmatched quantity shows its reason on the map.

7.2 Quantities, laws and paths to the new MAZE

- **By title.** A quantity whose Wikipedia title appears in a node's title list is placed under that node. This is the same title identity that links the new MAZE to the old. Nine quantities meet nodes this way: time, space, the Schwarzschild radius and the number of neurons are examples.
- **By dimension.** A rule exists to place a quantity by the unit stated in the paper, under the nodes for mass, space, time and charge (winding 4) and on the ladder (winding 12). The maps carry no units, so this rule places nothing yet. It is not bypassed with a dictionary of concepts.
- **By form.** Every law sits under `lus_helling`, because a law is a slope. Every path sits under `lus_overal`, because a path either closes or carries tension.

7.3 Papers as residents

The papers themselves become residents of the old MAZE: the 645 arXiv papers and the 1,303 PubMed records. They pass the same gate as any other item. The address comes from the content, the receipt names the arXiv identifier or PMID, and nothing is overwritten. Before any writing, 50 random papers were addressed without

storing; all 50 received distinct addresses (threshold 90%). On the maps, Δ behind a paper opens it at its address in the net, among its neighbours.

The result is a two-way thread. From a dot on a map the reader reaches the Wikipedia resident, its neighbourhood in meaning, and the derived node where one exists. From a paper the reader reaches its place in the net. The Transformer has become the entrance through which a research field enters both MAZEs.

8. Limitations

- **The assessor is a model.** All blind checks were done by a language-model agent. Planted errors and my own spot checks control for this only in part. A human expert in each field would be a stronger judge.
- **Language models in the loop.** Naming, law filtering and effect binding use a language model. Each step is cached, run at temperature 0 and tested blind. The model can still share blind spots with the assessor.
- **Rules derived on seen data.** The reading rules L13-L19 were derived from the same papers they now read. Only the test samples were fresh.
- **Physics paths untested.** No paper in the current corpora measures a chained pair directly.
- **Medicine paths are not causal.** They predict directions, not mechanisms, and they mix populations.
- **Name-only links.** A Wikipedia title can fit a name but not the field's meaning. Coverage is limited for compound quantities.
- **The theory.** The Vacuum.Net theory supplies the vocabulary: place, passage, trit, path, closure. The Transformer's results do not test the theory's physical claims. They test whether that vocabulary organises published claims in a useful way.

9. Reproducibility

All code, plans, blind samples, keys and outcomes are in the SWARP repository under `maze/transformer`. The main files are `v1.1/lees.py` and `kern.py` (Layer A), `layerB-v1.4/layerb.js` and `ai/` (Layer B), `koppeling.mjs` and `artikelen-opname.mjs` (the MAZE links). Each plan has a `.sha256` file committed before its run, and each outcome has an `uitkomst-*.md` file. The live API is documented at `maze.swarp.nl/handleidingen/transformer`. Its main routes are `/api/transformer/kaart`, `/api/transformer/geneeskunde` and `/api/transformer/koppeling?kaart=`.

Annotated references

1. Konstapel, J. (2026). *The Vacuum.Net Theory: Foundational Paper*, 8th ed., v8.1. Constable Research, Leiden. constable.blog.
The canonical statement of the theory: axioms N1-N5, theorems T1-T23, correspondences, predictions and open points, with a register of checks. Read it for the status of each claim.
2. Konstapel, J. (2026). *Curves in the Vacuum.Net*, v1.0. Constable Research, Leiden.

Derives paths, connection points and Theorem 8 (a power law is a straight line in position space). This is the direct theoretical basis of the Transformer.

3. Spencer-Brown, G. (1969). *Laws of Form*. London: Allen & Unwin.

The calculus of distinction from which the theory starts. Read it for the idea that one act of distinction, re-entering itself, generates a logic.

4. Kauffman, L. H. (1987). Self-reference and recursive forms. *Journal of Social and Biological Structures*, 10(1), 53–72.

Connects re-entry, eigenforms and imaginary values. It underlies the theory's step from distinction to oscillation.

5. Rowlands, P. (2007). *Zero to Infinity: The Foundations of Physics*. Singapore: World Scientific.

The nilpotent form of the Dirac equation that the Foundational Paper adopts. Read it to see where the theory takes Dirac as input.

6. Larson, D. B. (1959). *The Structure of the Physical Universe*. Portland: North Pacific.

The space-time dimensions used for the theory's units. A historical, unorthodox source; read it critically.

7. Knuth, D. E. (1997). *The Art of Computer Programming, Vol. 2: Seminumerical Algorithms*, 3rd ed., §4.1. Reading, MA: Addison-Wesley.

The standard treatment of balanced ternary, the number system behind every MAZE address.

8. Hayes, B. (2001). Third base. *American Scientist*, 89(6), 490–494.

A readable account of why base three is the most economical radix. It explains the trit as an engineering choice.

9. Barenblatt, G. I. (1996). *Scaling, Self-Similarity, and Intermediate Asymptotics*. Cambridge University Press.

Why power laws are the natural form of quantitative claims in physics. Background for Layer A.

10. West, G. B., Brown, J. H., & Enquist, B. J. (1997). A general model for the origin of allometric scaling laws in biology. *Science*, 276(5309), 122–126.

The classic example of power laws in biology. It explains why the life-sciences map exists at all.

11. Swanson, D. R. (1986). Fish oil, Raynaud's syndrome, and undiscovered public knowledge. *Perspectives in Biology and Medicine*, 30(1), 7–18.

The origin of literature-based discovery: $A \rightarrow B$ in one literature and $B \rightarrow C$ in another imply $A \rightarrow C$. The Transformer's paths are a quantitative version of this idea.

12. Smalheiser, N. R., & Swanson, D. R. (1998). Using ARROWSMITH: a computer-assisted approach to formulating and assessing scientific hypotheses. *Computer Methods and Programs in Biomedicine*, 57(3), 149–153.

The software successor to Swanson's method. Compare its word-based links with the Transformer's exponent- and effect-based links.

13. Tshitoyan, V., et al. (2019). Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*, 571, 95–98.

Shows that relations hidden across many papers can predict later findings. A modern counterpart that works with embeddings instead of explicit claims.

14. Fortunato, S., et al. (2018). Science of science. *Science*, 359(6379), eaao0185.

An overview of quantitative study of the scientific literature. It places the Transformer within that field.

15. DerSimonian, R., & Laird, N. (1986). Meta-analysis in clinical trials. *Controlled Clinical Trials*, 7(3), 177–188.
The random-effects method for pooling effect sizes. Background for how Layer B combines effects on one edge.
16. Higgins, J. P. T., & Thompson, S. G. (2002). Quantifying heterogeneity in a meta-analysis. *Statistics in Medicine*, 21(11), 1539–1558.
Defines I^2 , which the medicine map shows for edges with more than one effect.
17. Davey Smith, G., & Ebrahim, S. (2003). 'Mendelian randomization': can genetic epidemiology contribute to understanding environmental determinants of disease? *International Journal of Epidemiology*, 32(1), 1–22.
The design that could give medicine edges a causal direction, the planned next step for Layer B.
18. Lipscomb, C. E. (2000). Medical Subject Headings (MeSH). *Bulletin of the Medical Library Association*, 88(3), 265–266.
A short history of the vocabulary that supplies Layer B's places.
19. Lehmann, J., et al. (2015). DBpedia: a large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web*, 6(2), 167–195.
The source of the article summaries that the old MAZE uses to rank residents by meaning.
20. arXiv (2026). *arXiv API User's Manual* and bulk source access (arxiv.org/e-print). NCBI (2026). *Entrez Programming Utilities*.
The two public interfaces through which all papers in the Transformer were retrieved.

Code, plans and data: SWARP repository, [maze/transformer](https://github.com/maze/swarp). The maps: maze.swarp.nl/transformer. This essay as PDF: maze.swarp.nl/transformer/essay.pdf.