title: "The Design Step: Image Schemas, Theory Morphisms and a Provable Mythic Hemisphere"

J. Konstapel — Constable Research, Leiden — 22 September 2026

Sequel to "Clerk Maxwell and the Vacuum.Net Theory", sixth edition (21 September 2026). First edition of this paper. Status: written, not yet measured.

The next step of MAZE.AI is design: the passage between a rule and an image, in both directions. This paper makes that step as provable as the address kernel. The route is George Lakoff's. Abstract structure rests on a small set of image schemas, and a metaphor carries the structure of its schema into a new domain. The embodiment community has already given this a precise mathematical form. A metaphor is a theory morphism. A new concept is a colimit. An image schema is a small family of theories.

Three results follow for the machine. First, the gate can check an image the way it checks an address: axiom by axiom, with a named reason for every refusal. Second, the three positions of every schema are a trit, so the check is a finite table, and the mirror of a valid image is a valid image. Third, the boundary between machine and person falls exactly where the community itself drew it: the morphism is proved, the choice among images is attested.

The fixed image stays the fishnet. A schema is a mesh with three knots. A metaphor lays that mesh onto another part of the net and checks that every strand still closes.

Everything needed to follow every step is in the text. No reference has to be fetched.

1. Where the sixth edition ended, and an erratum

The sixth edition derived a machine with two hemispheres and one boundary. The address hemisphere (LAI) is a claim machine on frozen topology. It runs at maze.swarp.nl/ai over 8.19 million addressed pieces, with its kernel theorems machine-checked in Lean 4. The boundary is one gate with one output path. The tension hemisphere (RAI) received its operation, the ring measure R_Q, which passed its test at AUC 0.70 [0.61, 0.79] and now stands in the gate as tiebreaker.

One thing the sixth edition did not give the machine: a way to produce something that is not yet in the store. It can read what is there. It can name a vacancy. It cannot yet propose what would fill it. That missing operation is design, and it is the subject here.

Erratum to Section 16

The worked example of Section 16 is correct. Recomputed: $D_1 = (0.707, 0, 0, 0.707)$, $D_2 = (0.653, -0.653, -0.271, 0.271)$, R_Q = 0.827, naive mean 0.909, and the sign self-test leaves R_Q unchanged.

One sentence is wrong: "Flip the sense of the second cell ($u_2 \to (1, 0, -1)/\sqrt{2}$) … and R_Q drops toward 1/2." It does not drop. With $u_2 = (1, 0, -1)/\sqrt{2}$ the second winding becomes $D_2 = (0.653, 0.653, 0.271, 0.271)$. The x- and y-components change sign; the scalar part and the z-component do not.

The reason is a symmetry. The flip is the mirror $z \to -z$. The question $q = (1, 0, 0)$ and the answer $a = (0, 1, 0)$ lie in the mirror plane, and so does the first cell's axis. A rotation axis is a pseudovector, so under this mirror (v_x, v_y, v_z) becomes $(-v_x, -v_y, v_z)$. D_1 has only a scalar part and a z-component. The overlap stays $D_1 \cdot D_2 = 0.707 \cdot 0.653 + 0.707 \cdot 0.271 = 0.653$, and R_Q stays 0.827.

The replacement sentence: take a second cell that leans away from the answer, $u_2 = (0, -1, 1)/\sqrt{2}$. Then $D_2 = (0.271, -0.653, 0.271, -0.653)$. The overlap is $D_1 \cdot D_2 = 0.707 \cdot 0.271 - 0.707 \cdot 0.653 = -0.270$. With the sign-invariant form, $\lambda_1 = (1 + 0.270)/2 = 0.635$.

Same pair, same cosine (zero), R_Q 0.83 against 0.64. The cosine cannot see the difference; the ring measure is built from it. The seventh edition of the Maxwell paper carries this correction.

2. Design as Unity ↔ Mythic

Design is the passage between rule and image, in both directions. Will McWhinney's Paths of Change gives the frame. It distinguishes four realities. Unitary holds rules, principles and theory; its function is thinking. Sensory holds objects, facts and data. Social holds values and what matters. Mythic holds vision, symbols and ideas; its function is intuition.

Unitary and Mythic share one pole: both are monistic. They differ in one respect. Unitary is deterministic, Mythic volitional. Their combination is the Assertive mode of change, and McWhinney's example of it is the architect who designs a new building. Design is this pair.

The passage runs both ways, and each direction has its own work.

- **Unity → Mythic.** From what is fixed to an image of what is not yet there. The starting point already exists in the maze: a vacancy is a readable brief. The rules around an empty address say what shape would fit it.
- **Mythic → Unity.** From image to rule. The image is read back into addresses and axioms. Only then can the gate handle it.

The architect sees the building first and draws the structure afterwards. The image precedes the rule; the rule makes the image buildable.

The four realities place the rest of the machine. Unitary is the kernel and its proofs. Sensory is the store: pieces as material. Mythic is the field, still to be completed. Social is not a module at all. It is the net between nodes, where what matters is what returns in phase across many of them. A human attestation is where the loop touches the world.

This paper builds only the passage between Unitary and Mythic. It needs one new object, the image, and one new check at the gate.

3. Lakoff: image schema and invariance

Abstract structure rests on a small number of image schemas. That is the finding of George Lakoff and Mark Johnson, stated in 1980 and worked out in 1987. An image schema is a recurring pattern of bodily experience with an internal logic of its own. The standard ones are few: CONTAINER, PATH, VERTICALITY, CENTRE–PERIPHERY, PART–WHOLE, LINK and BALANCE.

Each schema carries inferences that come with it. Consider CONTAINER. If x is in y and y is in z, then x is in z. Nothing is in itself. If x is in y, then y is not in x. No one learns these as rules. They come with having a body that goes in and out of rooms.

A conceptual metaphor projects such a schema from a source domain onto a target domain. CATEGORIES ARE CONTAINERS projects the container schema onto classification. LIFE IS A JOURNEY projects the path schema onto a life. MORE IS UP projects verticality onto quantity.

The decisive part is the Invariance Principle (Lakoff 1990, 1993). A metaphorical mapping preserves the image-schema structure of the source, in a way consistent with the target. What was inside stays inside. What came later on the path comes later in the argument. The inferences of the source travel with the mapping.

Lakoff and Rafael Núñez applied this to mathematics itself (2000). Arithmetic, they argue, grows out of four grounding metaphors: collecting objects, building objects, measuring with a stick and moving along a path. The schemas are the floor; mathematics is built on them.

For MAZE.AI this is exactly the missing bridge. A vacancy has no content, but it has a shape: the schema its neighbours imply. An image is a schema filled from another domain. The Invariance Principle says when the filling is legitimate. What remains is to make that principle checkable. The next section does so.

4. The mathematics: theory, morphism, colimit

A metaphor that satisfies the Invariance Principle is a theory morphism, and a blend of two concepts is a colimit. This is the form Joseph Goguen gave to the work of Lakoff, Fauconnier and Turner, inside the institution theory he built with Rod Burstall. Four definitions and one lemma carry everything that follows.

Definitions

Signature. A signature Σ is a set of sorts and a set of predicate and function symbols, each with its sorts fixed. Example: one sort *Thing*, one binary predicate *In*.

Theory. A theory $T = (\Sigma, A)$ is a signature with a set A of axioms, which are sentences over Σ. A model of T interprets the sorts as sets and the symbols as relations and functions, such that every axiom holds.

Signature morphism. A map $\sigma : \Sigma \to \Sigma'$ sends sorts to sorts and symbols to symbols, respecting sorts. It translates every sentence ϕ over Σ into a sentence $\sigma(\phi)$ over Σ', symbol by symbol. It also turns every model M' of Σ' into a model $M'|_\sigma$ of Σ, its reduct: read each source symbol as the target symbol it maps to. The two directions are tied by the satisfaction condition:

$$M' \models \sigma(\varphi) \iff M'|_{\sigma} \models \varphi$$

Theory morphism. A signature morphism σ is a theory morphism from (Σ, A) to (Σ', A') when every translated axiom follows in the target:

$$A' \vdash \sigma(a) \quad \text{for every } a \in A$$

This is the checkable form of "the mapping is consistent with the target". It is a finite list of obligations, one per source axiom.

Lemma 1 (invariance)

If σ is a theory morphism and ϕ follows from the source axioms, then $\sigma(\phi)$ follows from the target axioms.

$$A \vdash \varphi ;\Longrightarrow; A' \vdash \sigma(\varphi)$$

Proof. Let M' be any model of A'. By the morphism condition, M' satisfies $\sigma(a)$ for every source axiom a. By the satisfaction condition, the reduct $M'|\sigma$ satisfies every a, so it is a model of A. Since A entails ϕ, the reduct satisfies ϕ. By the satisfaction condition again, M' satisfies $\sigma(\phi)$. This holds for every model of A', so A' entails $\sigma(\phi)$. ∎

Lemma 1 is Lakoff's Invariance Principle as a theorem. Every inference of the source schema travels to the target, and the gate has to check only the axioms, never the inferences.

Worked example: a morphism, and a refusal

Source: CONTAINER. Sort *Thing*, predicate *In*. Two axioms:

- C1 (transitivity): if In(x, y) and In(y, z), then In(x, z).
- C2 (irreflexivity): not In(x, x).

One inference follows. C3 (asymmetry): if In(x, y), then not In(y, x). *Proof:* suppose In(x, y) and In(y, x). C1 with z = x gives In(x, x), which C2 forbids. ∎

Target: CLASSIFICATION. Sort *Concept*, predicate *Sub* ("is a strict subkind of"). Axioms K1: Sub is transitive. K2: nothing is a strict subkind of itself.

The morphism σ sends *Thing* to *Concept* and *In* to *Sub*. The obligations are two. σ(C1) is K1; σ(C2) is K2. Both hold, so σ is a theory morphism: CATEGORIES ARE CONTAINERS passes.

By Lemma 1 the target gains σ(C3) for free: if dog is a strict subkind of mammal, mammal is not a strict subkind of dog. Nobody wrote that axiom. It travelled with the schema. Two checks, one theorem gained.

A refusal: FRIENDSHIP. Sort *Person*, predicate *Friend*. Axioms: Friend is symmetric and irreflexive. Try the same σ, with *In* sent to *Friend*. The first obligation fails: σ(C1) is transitivity of friendship, and a model refutes it. Take three people a, b, c with a–b and b–c friends and a–c not. The gate refuses and names C1. It is worse than a gap: σ(C3) would make friendship asymmetric, which contradicts the target outright.

Blend as colimit

A blend combines two input theories over a shared generic theory G. Given morphisms $G \to T_1$ and $G \to T_2$, the colimit (here a pushout) is the smallest theory B that receives both inputs and agrees on G. Its signature is the union of the two, with the symbols that come from G identified. Its axioms are the union of the translated axioms.

Worked blend. G has one sort X, one predicate R and one axiom: R is transitive. T_1 is PATH: sort *Stage*, predicate *Before*, transitive, irreflexive and total (any two distinct stages are ordered). T_2 is VERTICALITY: sort *Level*, predicate *Above*, transitive and irreflexive. Map R to *Before* and to *Above*. The pushout identifies *Stage* with *Level* and *Before* with *Above*. Its axioms: transitive, irreflexive, total.

The blend is a path that climbs. It also carries something neither input stated about height: any two levels are comparable, because totality came in from PATH. That is the structure behind "getting ahead" and "climbing the ladder". The new content of a blend is exactly the axioms one input brings to the other.

Two facts limit the blend, and both were stated by the community itself. A colimit can be inconsistent; the COINVENT project repaired such blends by dropping axioms until they were consistent. And there are usually many possible blends; Goguen and Harrell state that the optimality principles that select among them cannot be captured formally. Section 6 turns these two facts into the design of the gate.

5. The seven schemas in trits, and the mirror lemma

Every standard image schema has three positions, and they are a trit. CONTAINER has inside, boundary and outside. PATH has goal, route and source. VERTICALITY has up, level and down. The middle position is the passage between the two ends. The same register that writes addresses in the address hemisphere writes schema positions here.

The trit schema

A trit schema S is a theory with three ingredients:

1. a sort E of entities and the fixed sort Trit = $\{-1, 0, +1\}$, shared by every schema and never renamed;
2. three role constants r_+, r_0, r_-, one per position;
3. the schema's own relation and at most three axioms.

A morphism between trit schemas, or from a schema into a domain theory, is **trit-preserving** when it leaves Trit fixed and sends each role to a role of the same position. That is three equalities, checkable by table.

One axiom recurs, and it is the trit's own. **Passage:** an entity that is at -1 at time t and at $+1$ at a later time t' is at 0 at some moment in between. Nothing jumps from one end to the other. In balanced ternary this is the statement that each step changes a digit by one.

The seven

Schema	1	0	-1	Axioms
CONTAINER	insid	bounda	outsid	In is transitive and irreflexive; passage
PATH	goal	route	sourc	Before is transitive, irreflexive and total; passage
VERTICALITY	up	level	down	Above is transitive and irreflexive; passage; unforced motion runs from +1 toward −1
CENTRE–PERIPHERY	centr e	betwee n	periph ery	Closer is transitive and irreflexive; passage
PART–WHOLE	whol	connec	part	PartOf is transitive and irreflexive; a whole without one of
LINK	A	bond	B	Linked is symmetric; what moves one end moves the other
BALANCE	one side	equilib rium	other side	at equilibrium, $w(+1) + w(-1) = 0$

BALANCE deserves a note. Its axiom is the closure law of the net itself: a closed winding carries net zero. The schema through which people feel equilibrium is the same law the address register obeys.

The gate's obligations for one image are therefore small and fixed: three role equalities plus at most three axioms. At most six checks per image, whatever the domain.

Lemma 2 (mirror)

Let μ be the mirror: $+1 \mapsto -1$, $-1 \mapsto +1$, $0 \mapsto 0$, with each ordered relation sent to its converse. Then the seven schemas are closed under μ, and LINK and BALANCE are fixed by it.

Proof. Transitivity, irreflexivity and totality of a relation hold exactly when they hold for its converse. Passage names the two ends symmetrically, so swapping them leaves it unchanged. The direction axiom of VERTICALITY becomes "unforced motion runs from −1 toward +1" for the converse relation *Below*, which is VERTICALITY read from the other end. LINK is symmetric, so the converse of *Linked* is *Linked*. BALANCE reads $w(+1) + w(-1) = 0$, and swapping the sides gives $w(-1) + w(+1) = 0$, the same equation. ∎

Corollary. If $\sigma : S \to T$ is a trit-preserving theory morphism, then σ composed with μ is a theory morphism from the mirrored schema into the same target. Every valid image comes with a valid mirror reading: inside read as outside, source as goal, up as down. For LINK and BALANCE the image is its own mirror.

This is the mirror-address corollary of the foundational paper, one level up. Every address has a sign-mirrored partner, and the zero address is its own mirror. Every image has a mirror image, and the balanced schemas are their own. In practice the corollary doubles the reach of every admitted image: the mirror reading is a candidate for the mirrored vacancy, at no extra proof.

6. The gate for images: what is proved and what is attested

The gate checks an image the way it checks an address: by finite obligations, each passed or refused with a name. Nothing about the gate's contract changes. One output path, no second one. The field proposes, the gate checks, the person attests.

What an image is

An image for a vacancy v consists of five parts:

1. a schema S from Section 5;
2. a source theory T_s: a domain where S is fully inhabited, with its axioms;
3. a trit-preserving morphism $\sigma : S \to T_s$, showing the source really carries the schema;
4. the target theory T_v: what the store holds around the vacancy, as axioms, with one symbol marked open — the symbol the vacancy lacks;
5. a map τ from the source to the target, on the symbols they share.

In the fishnet: the schema is a mesh with three knots. The source is a place in the net where that mesh is tied fully. The target is the torn place. τ lays the tied mesh over the tear.

The three checks

Check 1 — the source carries the schema. σ must be a trit-preserving theory morphism. Three role equalities, at most three axioms.

Check 2 — the mapping keeps its structure. τ must be a theory morphism on the shared axioms: every source axiom that τ translates must follow in the target. By Lemma 1, every inference of the source then holds in the target too.

Check 3 — the image does not presuppose the vacancy. The vacancy names two symbols: the open one, and the target symbol the open one has to yield. No primitive of the source may map onto that yielded symbol. An image that already uses the result to produce the result fills nothing. The check is purely syntactic: compare symbols.

Each check that fails refuses the image and names the obligation that failed: the axiom that did not follow, or the symbol that closed the circle. A refusal is a result, booked in the ledger like any other.

What the image adds

An image that passes brings with it the source axioms that have no counterpart in the target yet. Those are the new content, exactly as in the blend of Section 4. The gate does not check them against the world. It marks them *proposed*: the design the image carries, awaiting attestation.

The boundary

Proved by the machine	Attested by the person
The source carries the schema (check 1)	Whether the image is good
The mapping preserves every inference of the source (check 2, Lemma 1)	Which of several passing images to keep
The image does not presuppose the vacancy (check 3)	Whether the proposed new axioms hold in the world
The mirror reading is valid as well (Lemma 2)	What the image means

The right-hand column is where Goguen and Harrell drew the limit of formalisation: meaning and the optimality of a blend. It is where the sixth edition drew it too: the store is proved, the settlement is human. The two boundaries are one boundary.

The claim of the gate is bounded accordingly. It claims one thing: this map is a theory morphism on these axioms. It never claims that the image is true, or that it is the best one.

In the machine

The proofs live in the Lean kernel, beside the address theorems. The seven schemas become Lean definitions. Lemma 1 and Lemma 2 become Lean theorems, proved once. Each image produces its obligations as Lean goals: at most six for check 1, one per shared axiom for check 2. Mathlib's model-theory library offers first-order languages, maps between them and their action on sentences; whether it covers the satisfaction condition in the form needed here is for Claude Code to establish. Where the target axioms are numerical, the obligations are closed by computation.

The field's part is to propose. Given a vacancy, it searches the bank for sources where the vacancy's schema is fully inhabited. The ring measure R_Q can order the candidates before the gate sees them. The ordering is advice; the gate's checks are the proof; the person's attestation is the settlement.

7. Worked case: O1, Bjerknes and the rubber sheet

The first real vacancy passes the gate with one image and refuses another, each for a named reason. The vacancy is O1 from Finding 5 of the sixth edition: the note on gravitation in Maxwell's 1865 paper.

The vacancy

Maxwell computed the intrinsic energy of the gravitational medium and found it lower wherever a resultant force exists. Then he stopped. He could not see in what way a medium could possess such properties. He had the rule and no image. That is a vacancy in the exact sense of Section 6.

The target theory T_v, per unit volume of the medium:

- M1: $e = C - R^2/8\pi$, with e the energy density, C the energy of the undisturbed medium and R the resultant force.
- M2: $e \geq 0$ everywhere, so $C \geq R^2/8\pi$ for every force that occurs. C must be enormous.
- M3: two bodies attract, with a force proportional to $1/d^2$.
- M4: bodies move toward configurations of lower total energy.

Open symbol: *Carrier*, the property of the medium by which e is lowered where R exists. Yielded symbol: *Attract*.

The schema is VERTICALITY. +1 is the base level C, 0 is the passage carried by R, −1 is the lowered energy.

Image 1: pulsating spheres (Bjerknes)

Carl Anton Bjerknes and his son Vilhelm studied spheres that pulsate in an ideal fluid. The source theory T_s, with pressure time-averaged:

- B1 (Bernoulli): $p = p_0 - (\varrho/2)\cdot u^2$, with p_0 the ambient pressure, ϱ the density and u the flow speed.
- B2: $p \geq 0$ (no cavitation), so $p_0 \geq (\varrho/2)\cdot u^2$.
- B3: two spheres pulsating in phase attract; out of phase they repel. The mean force between them is $F = \varrho\cdot\langle\dot{V}_1\dot{V}_2\rangle / (4\pi\, d^2)$, with V_1, V_2 the sphere volumes and d their distance.
- B4: bodies are pushed from higher mean pressure toward lower.

Check 1. σ sends up to p_0, level to the flowing region, down to the lowered pressure. By B1 the pressure is below p_0 wherever $u \neq 0$, so the order holds. Pressure varies continuously, so passage holds. B4 is the direction axiom. Passed: three role equalities, three axioms.

Check 2. τ sends p to e, p_0 to C, u to R and the constant $\varrho/2$ to $1/8\pi$. The first obligation is the most telling. Set the two laws side by side:

$$p = p_0 - \tfrac{\rho}{2},u^2 \qquad\longmapsto\qquad e = C - \tfrac{1}{8\pi},R^2$$

τ(B1) is M1, symbol for symbol. τ(B2) is M2. τ(B4) is M4. The distance law of B3 goes to the distance law of M3: $1/d^2$ to $1/d^2$. Four obligations, four passed. By Lemma 1, every inference of the pulsating-sphere theory about the energy of the medium now holds for Maxwell's medium too.

Check 3. The primitives of the source are fluid, pressure, flow and pulsation. None of them maps onto *Attract*. Attraction is what the source produces, not what it assumes. Passed.

What the image adds. One axiom of the source has no counterpart in the target: B3's phase condition. Carried over, it reads: two bodies attract gravitationally when their pulsations are in phase, and would repel if they were out of phase. Gravity shows no repulsion. The image therefore proposes that all matter pulsates in one common phase.

The gate marks this *proposed*, not proved. In the language of the net it is recognisable at once. What persists is what returns in phase; the ambient pressure p_0 is the enormous base tension C of the free net; the order parameter $r \approx 1$ of the sixth edition is the global state the proposal asks for. Whether the proposal holds in the world is for attestation.

The mirror reading. By Lemma 2 the mirrored image is valid as well: raised pressure between the bodies, pushing them apart. That is B3's out-of-phase half. The mirror names the mirrored vacancy in the target: why gravity has no repulsive branch. The image answers it in the same stroke: no body is out of phase.

Image 2: the rubber sheet

The familiar picture of gravity is a ball on a stretched rubber sheet that dips under a heavy mass. The source theory:

- RS1: a sheet under tension T with a point load F dips as $h(r) = (F / 2\pi T)\cdot\ln(r / r_0)$, for small deflection.
- RS2: a ball on the sheet rolls down the slope because of *Pull*, the downward force on the ball.
- RS3: the force on the ball follows the slope, $dh/dr = F / (2\pi T\cdot r)$, so it falls off as $1/r$.

Check 1 passes: flat sheet up, slope as passage, dip down.

Check 2 fails at RS3. τ sends the 1/r law of the sheet to a 1/r law of attraction. M3 demands $1/d^2$. A logarithmic dip gives the wrong distance law, so τ(RS3) does not follow; it contradicts the target. Refused, with RS3 named.

Check 3 fails at RS2. *Pull* is a primitive of the source. The only target symbol it can map to is *Attract*, the yielded symbol. The picture uses the Earth's gravity to explain gravity. Refused, with *Pull* named.

The record

Image	Check 1	Check 2	Check 3	Outcome
Bjerknes, pulsating	passed	4 of 4	passed	admitted as morphism; phase condition proposed, awaiting attestation
Rubber sheet	passed	failed at RS3 (1/r against $1/d^2$)	failed at RS2 (*Pull* ↦ *Attract*)	refused, two reasons booked

A first run of this case was made by hand in a small SQLite prototype of the bank. There the facts per source were entered as flags, and the distance test was added for the occasion. The theory form of this section replaces both: the facts become axioms, and the distance test becomes an ordinary obligation of check 2.

8. The bank: five tables

The design step needs one new store beside the pieces: a bank of schemas, metaphors and images. Five tables carry it. The schemas are grammar and live with the kernel. The metaphors and images are inhabitants of the maze: each gets an address and passes the same gate as any other piece.

Table	One row is	Fields	Filled by
sche	one trit schema	id, name, role +1, role 0, role	the seven of Section 5, once, with the Lean
metaphors	one conceptual metaphor	id, name, source domain, target domain, schema, source	parsed from published catalogues, each schema assignment attested
roles	one role of a metaphor	metaphor, source role, target role, position (+1/0/−1),	the existing coder, per role
examples	one sentence where the	metaphor, piece address, sentence	a search of the existing corpus for the source-domain words
images	one proposal for a vacancy	vacancy, metaphor, filling (role → address), obligations,	the field proposes; the gate writes the results; status is proposed, admitted as

The first filling does not start from nothing. The Master Metaphor List (Lakoff, Espenson and Schwartz, Berkeley 1991) catalogues conceptual metaphors in the form X IS Y, with their sub-mappings. MetaNet (ICSI Berkeley) organises metaphors as a network of schemas, frames and mappings, which is the closest existing structure to what the bank needs. The Mapping Metaphor project (Glasgow, 2015) charts metaphorical links across the history of English. The VU Amsterdam Metaphor Corpus (Steen and colleagues, 2010) marks metaphor in running text by an explicit procedure.

For the formal side, the formalised image schemas of Hedblom, Kutz and Neuhaus, and the CASL specifications of the COINVENT project, are ready sources of axioms. What is available under which licence is for Claude Code to establish before any import.

The corpus already in the maze — the encyclopaedia, the art and music collections, the blog — contains the metaphors in use. The examples table grows from it. A metaphor with many examples is a well-trodden strand; one with none is a candidate still waiting for use.

9. Status

This paper makes five kinds of statement, and keeps them separate.

Borrowed and established. Lakoff's image schemas and Invariance Principle; McWhinney's four realities and six modes; Goguen's reading of metaphor as morphism and blend as colimit; the institution framework of Goguen and Burstall; the formalised image schemas of Hedblom, Kutz and Neuhaus; the COINVENT blends in mathematics; Bernoulli's law and the Bjerknes force; Maxwell's 1865 note. These are cited, not claimed.

Proved here. Lemma 1 (invariance under a theory morphism), with its proof in Section 4. Lemma 2 (mirror), with its proof in Section 5. The worked morphism and refusal of Section 4. The erratum to Section 16 of the sixth edition, with its numbers. Every step is in the text.

Constructed here. The trit schema and the passage axiom; the three checks of the gate for images; the five-part image; the five tables. These are design choices anchored in the theory, not theorems.

Proposed, awaiting attestation. The phase reading of gravitation that the Bjerknes image carries (Section 7): all matter pulsating in one common phase. The gate admitted the morphism. It did not, and cannot, admit the proposal.

Not yet built, not yet measured. The work lies with Claude Code. It writes the seven schemas and both lemmas in the Lean kernel. It creates the five tables and fills them from the catalogues whose licences allow it. It runs the gate for images over the current vacancy list. Per vacancy it reports three counts: images admitted as morphisms, images refused with their named obligations, and images refused as circular. It books all outcomes in the ledger. The owner attests the admitted proposals.

What this edition leaves open is therefore not writing. It is the Lean formalisation, the first run over many vacancies, and the attestations.

10. Annotated references

J. Konstapel, "Clerk Maxwell and the Vacuum.Net Theory", sixth edition, Constable Research, 21 September 2026. *Why read?* The paper this one continues: the two hemispheres, the gate, the ring measure R_Q and the open calculation O1. *Reading advice:* Sections 7–9 for the architecture, Section 16 beside the erratum above, Finding 5 for O1.

G. Lakoff and M. Johnson, *Metaphors We Live By*, University of Chicago Press, 1980. *Why read?* The founding statement that ordinary abstract thought is structured by metaphor. *Reading advice:* short and plain; the first ten chapters carry the argument.

M. Johnson, *The Body in the Mind*, University of Chicago Press, 1987. *Why read?* The image schema worked out from bodily experience, with the internal logic of CONTAINER, PATH, BALANCE and others. *Reading advice:* the chapter on the "in–out" orientation shows how inferences come with a schema.

G. Lakoff, *Women, Fire, and Dangerous Things*, University of Chicago Press, 1987. *Why read?* Categories as containers and the case studies behind the schema theory.

G. Lakoff, "The Invariance Hypothesis: is abstract reason based on image-schemas?", *Cognitive Linguistics* 1(1), 39–74, 1990. *Why read?* The first statement of the principle this paper turns into Lemma 1.

G. Lakoff, "The contemporary theory of metaphor", in A. Ortony (ed.), *Metaphor and Thought*, 2nd ed., Cambridge University Press, 1993. *Why read?* The Invariance Principle in its mature form, with the target-domain override. *Reading advice:* the best single paper to start with.

G. Lakoff and R. Núñez, *Where Mathematics Comes From*, Basic Books, 2000. *Why read?* Arithmetic and beyond built from four grounding metaphors: the embodiment programme applied to mathematics itself.

W. McWhinney, *Paths of Change: Strategic Choices for Organizations and Society*, Sage, 1992 (rev. 1997). *Why read?* The four realities and six modes, including the Assertive mode that this paper reads as design.

J. Goguen and R. Burstall, "Institutions: abstract model theory for specification and programming", *Journal of the ACM* 39(1), 95–146, 1992. *Why read?* The satisfaction condition and theory morphisms in full generality: the mathematics under Section 4. *Reading advice:* the first sections suffice.

J. Goguen, "An introduction to algebraic semiotics, with application to user interface design", in C. Nehaniv (ed.), *Computation for Metaphors, Analogy, and Agents*, LNAI 1562, Springer, 1999. *Why read?* Sign systems as theories, semiotic morphisms, and conceptual blending as a colimit.

J. Goguen and D. F. Harrell, "Style: a computational and conceptual blending-based approach", in S. Argamon et al. (eds.), *The Structure of Style*, Springer, 2010. *Why read?* Where the authors state what cannot be formalised: meaning and the optimality of a blend. This is the right-hand column of Section 6.

G. Fauconnier and M. Turner, "Conceptual integration networks", *Cognitive Science* 22(2), 133–187, 1998. *Why read?* The four-space picture of blending that Goguen made exact.

W. Kuhn and A. U. Frank, "A formalization of metaphors and image-schemas in user interfaces", in D. Mark and A. Frank (eds.), *Cognitive and Linguistic Aspects of Geographic Space*, Kluwer, 1991. *Why read?* The first algebraic definitions of CONTAINER and SURFACE.

A. Rodríguez and M. Egenhofer, "Image-schemata-based spatial inferences: the container-surface algebra", COSIT 1997, LNCS 1329, Springer. *Why read?* Schemas that interact, with explicit operations and inferences.

A. U. Frank and M. Raubal, "Formal specification of image schemata — a step towards interoperability in geographic information systems", *Spatial Cognition and Computation* 1, 67–101, 1999. *Why read?* The survey of early formalisations, written in Haskell.

W. Kuhn, "An image-schematic account of spatial categories", COSIT 2007, LNCS 4736, Springer. *Why read?* An algebraic theory in which LINK, CONTAINER, SUPPORT and PATH carry the affordances of categories.

M. M. Hedblom, O. Kutz and F. Neuhaus, "Choosing the right path: image schema theory as a foundation for concept invention", *Journal of Artificial General Intelligence* 6(1), 2015. *Why read?* Image schemas as families of theories, used to steer blending. *Reading advice:* the PATH family is the model for Section 5.

M. M. Hedblom, *Image Schemas and Concept Invention*, Springer, 2020. *Why read?* The full treatment, with an image schema logic.

T. Mossakowski, C. Maeder and K. Lüttich, "The Heterogeneous Tool Set, Hets", TACAS 2007, LNCS 4424, Springer. *Why read?* The proof tool with which the community computes colimits and discharges morphism obligations; the counterpart of the Lean kernel here.

F. Bou, M. Schorlemmer, J. Corneli, D. Gómez-Ramírez, E. Maclean, A. Smaill and A. Pease, "The role of blending in mathematical invention", ICCC 2015. *Why read?* COINVENT's case that blending produces new mathematics, with worked blends.

D. A. J. Gómez-Ramírez, *Artificial Mathematical Intelligence*, Springer, 2020. *Why read?* Blending as colimits of many-sorted theories, with meta-theorems generating concepts of field and Galois theory, computed in Hets.

R. Confalonieri et al. (eds.), *Concept Invention*, Springer, 2018. *Why read?* The collected results of COINVENT, including the repair of inconsistent blends.

D. Gentner, "Structure-mapping: a theoretical framework for analogy", *Cognitive Science* 7(2), 155–170, 1983. *Why read?* The parallel tradition: analogy as mapping of relational structure, with systematicity as selection.

V. Bjerknes, *Fields of Force*, Columbia University Press, 1906. *Why read?* The hydrodynamic analogies of Carl Anton and Vilhelm Bjerknes, including the attraction of in-phase pulsating spheres used in Section 7.

J. C. Maxwell, "A Dynamical Theory of the Electromagnetic Field", *Philosophical Transactions of the Royal Society* 155, 459–512, 1865. *Why read?* Section 82, the note on gravitation: the formula of the vacancy and the sentence where Maxwell stopped.

The Mathlib Community, *Mathlib*, model theory library (FirstOrder.Language and maps between languages). *Why read?* The Lean library in which the kernel can state signatures, sentences and their translation.