From Data Infrastructure to Structural Intelligence

Learn, Preserve, Verify and Interpret for Trustworthy AI Agents



Prof. Dr. Shang-Ling Jui

Applied Mathematics & AI Algorithms

OLYMPIA AI Data Infrastructure Summit 2026
Heraklion, Crete, Greece | 4–5 November 2026

HUAWEI

AI-Guided Chip Design: The Acceptance Problem

Coupled structures, admissible search and independent acceptance checks

Learn — Learn from data, physics and prior designs

Preserve — Maintain multi-layer design constraints

Verify — Multi-physics, multi-tool, multi-level validation

Interpret — Engineering judgment and value creation

1 Problem Definition: Chip Design as Constrained Structure Generation

Design representation

$$S = (G, \theta, X, R, \cdots)$$

- G : circuit graph / netlist (discrete)
- θ : device and circuit parameters (continuous)
- X : geometry / placement (continuous)
- R : routing (discrete + continuous)
- $\cdots$: architecture, clock, power, etc.

Design objectives (multi-objective)

$$\min_{S \in A_{design}} \; J(S)$$

$$J(S) = \left[P(S), D(S), A(S), T(S), C(S), \cdots\right]^T$$

- P : power D : delay A : area
- T : temperature C : cost Y : yield
- $\cdots$: reliability, manufacturability, etc.

Admissible design space

$$\mathcal{A}_{design} = \bigcap_i \mathcal{A}_i$$

- $\mathcal{A}_{logic}$: logical correctness
- $\mathcal{A}_{electrical}$: electrical constraints
- $\mathcal{A}_{physical}$: geometric/physical constraints
- $\mathcal{A}_{manufacturing}$: process constraints
- $\cdots$: timing, power, thermal, reliability

2 Mathematical Structure of Design Constraints

Multiple mathematical formalisms define the admissible design space.

Logical & Combinatorial Constraints

$$G \in \mathcal{G}_{valid}, \quad G \models \Phi_{topology}$$

- Graph theory, combinatorial optimization
- SAT/SMT, MSO logic
- Connectivity, fan-in/out, design rules



Circuit Graph G

Geometric & Physical Constraints

$$X \in X_{valid}, \quad g_i(X) \le 0$$

- Non-overlap, density, symmetry
- DRC rules, layout constraints
- Topological and geometric invariants

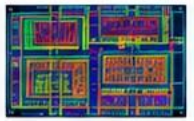

Layout X

Electrical & Timing Constraints

$$f_c(S) \le 0, \quad T_{path}(S) \le T_{max}$$

- Circuit laws (KCL/KVL)
- Timing analysis, signal integrity
- Power, noise, stability

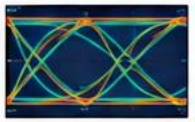

Timing / Signal

Physics & PDE Constraints (e.g., thermal, EM, photonic)

$$\mathcal{D}(u; \bar{\Omega}) = 0, \quad \Omega \in \Omega_{valid}$$

- Maxwell, heat, fluid, quantum
- PDE-constrained optimization
- Inverse design, topology optimization

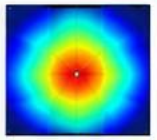

Thermal / EM Field

3 AI-Guided Design Generation and Search

Propose structured candidates.

Structure-aware Generative Models

$$p_\theta(S) = p(G, \theta, X, R, \cdots)$$

- Graph neural networks
- Transformer / diffusion models
- Physics-informed models
- Surrogate models
- Uncertainty estimation

Constraint-Preserving Search

$$S_{t+1} = T(S_t, a_t), \quad a_t \in A(S_t)$$

- Preserve hard process constraints
- Reinforcement learning / MCTS
- Heuristics + differentiable optimization
- Multi-scale and hierarchical search

AI-generated design candidates



Circuit Topology G | Layout X | Device Parameters θ

4 Multi-Level Verification (A Composition of Verifiers)

Checks under stated assumptions.

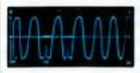

Functional Verification V_{func}
Simulation, formal verification

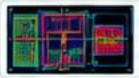

Physical Verification V_{DRC}
DRC, LVS, layout consistency



Timing Verification V_{timing}
STA, timing closure

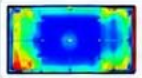

Power/Thermal Verification V_{power}
Power analysis, IR drop, thermal

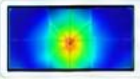

Physics Verification $V_{physics}$
EM/thermal/optical simulation



Manufacturing Verification V_{manu}
Process variation, yield, reliability

Overall verification

$V_i(S, \Phi_i; E_i) \in \{Pass, Fail, Unknown\}$
Accept: all mandatory checks pass
Same snapshot and specification

5 Interpretation and Engineering Decision

A verified design is not necessarily a good product.

Pareto Optimization and Trade-offs

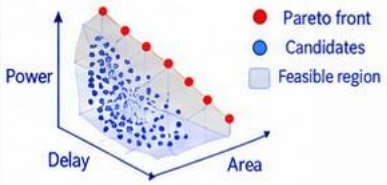

Human Engineer (with AI)



- Meet product requirements?
- Robust and manufacturable?
- Cost, yield, development time?
- Elegant and maintainable?
- Fit into roadmap?
- Better alternatives?

Engineering value creation

- Final design decision
- New design methodologies (Design for AI)
- Reusable IP and knowledge
- Continuous improvement

Feedback Loop: Simulation / Verification Results → Error Analysis → Model Update → Improved Search → Next Iteration

LPVI Fig. 7, locally revised. Search preserves process constraints. Final acceptance checks the target specification.

Huawei Proprietary - Restricted Distribution

HUAWEI

LPVI: Four Responsibilities for Structural Intelligence

Probability, structure, verification and meaning

LPVI Fig. 1, locally revised

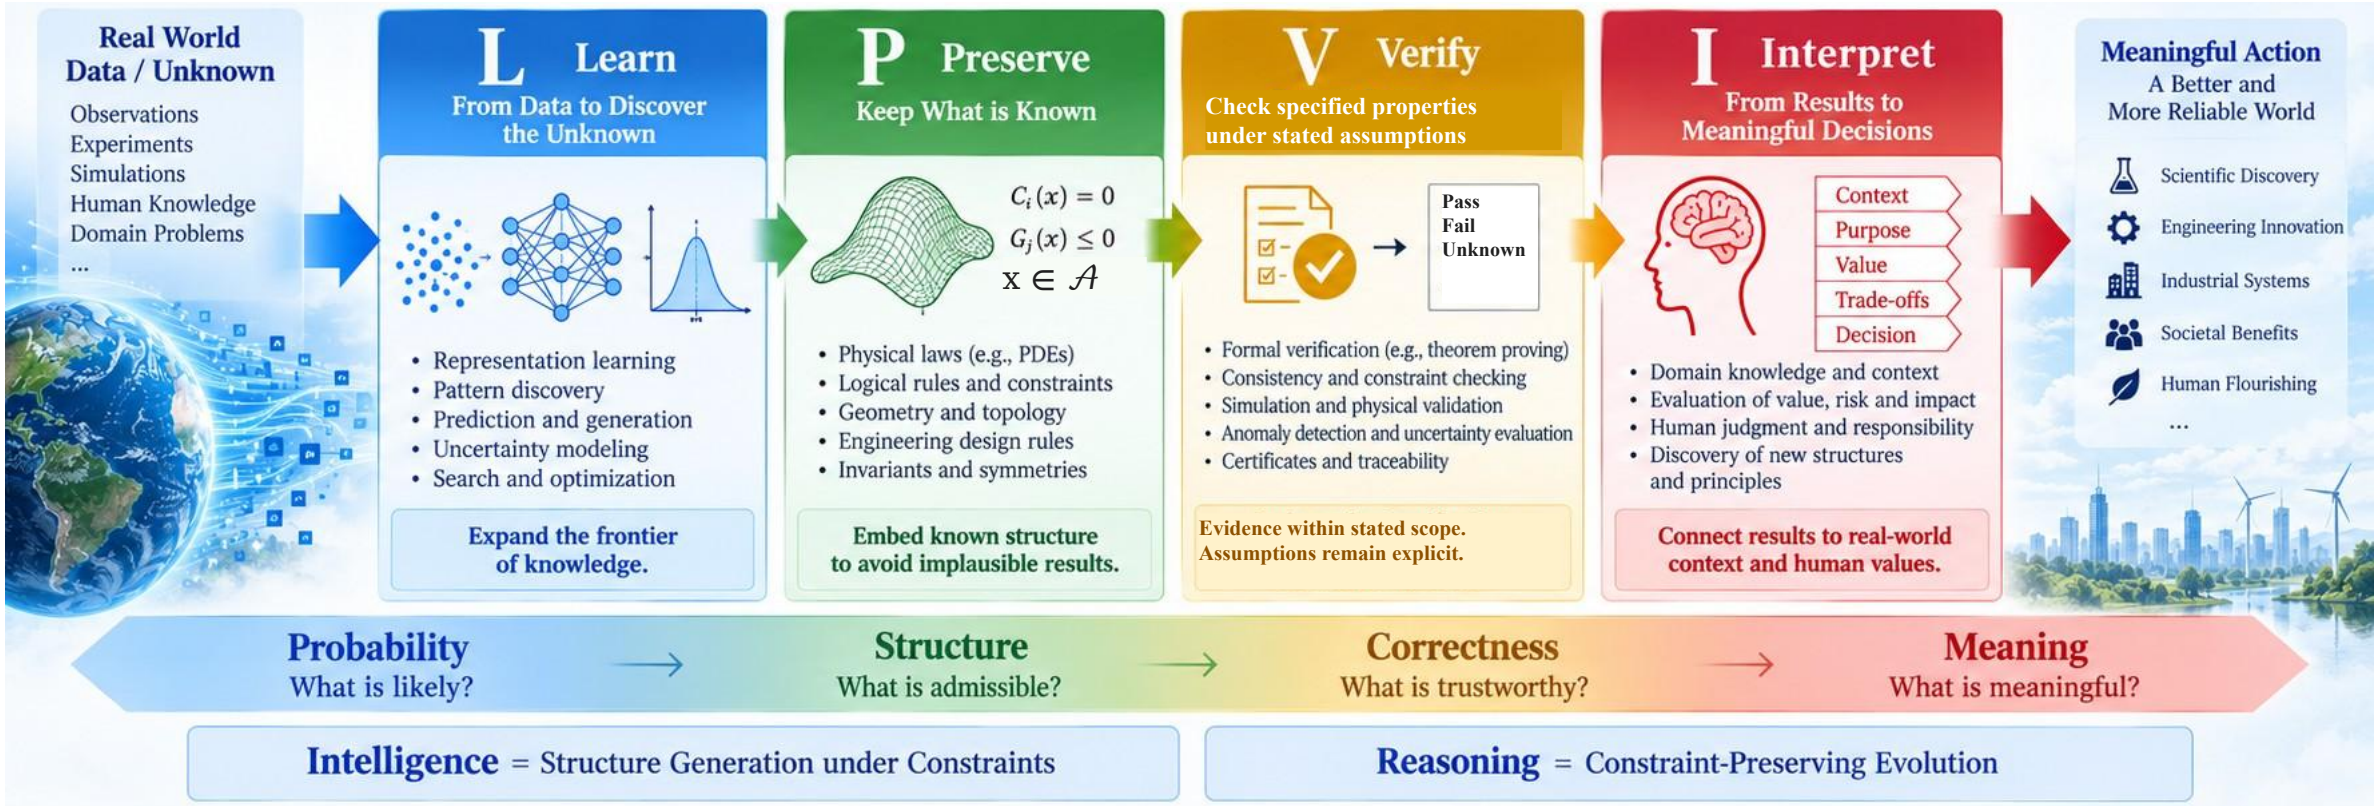

The arrows organize an iterative workflow.
Each responsibility has its own assumptions and evidence obligations.

HUAWEI

Constraint-Preserving Evolution

Forward invariance under a fixed deterministic operator

State Space, Admissible Set and Evolution

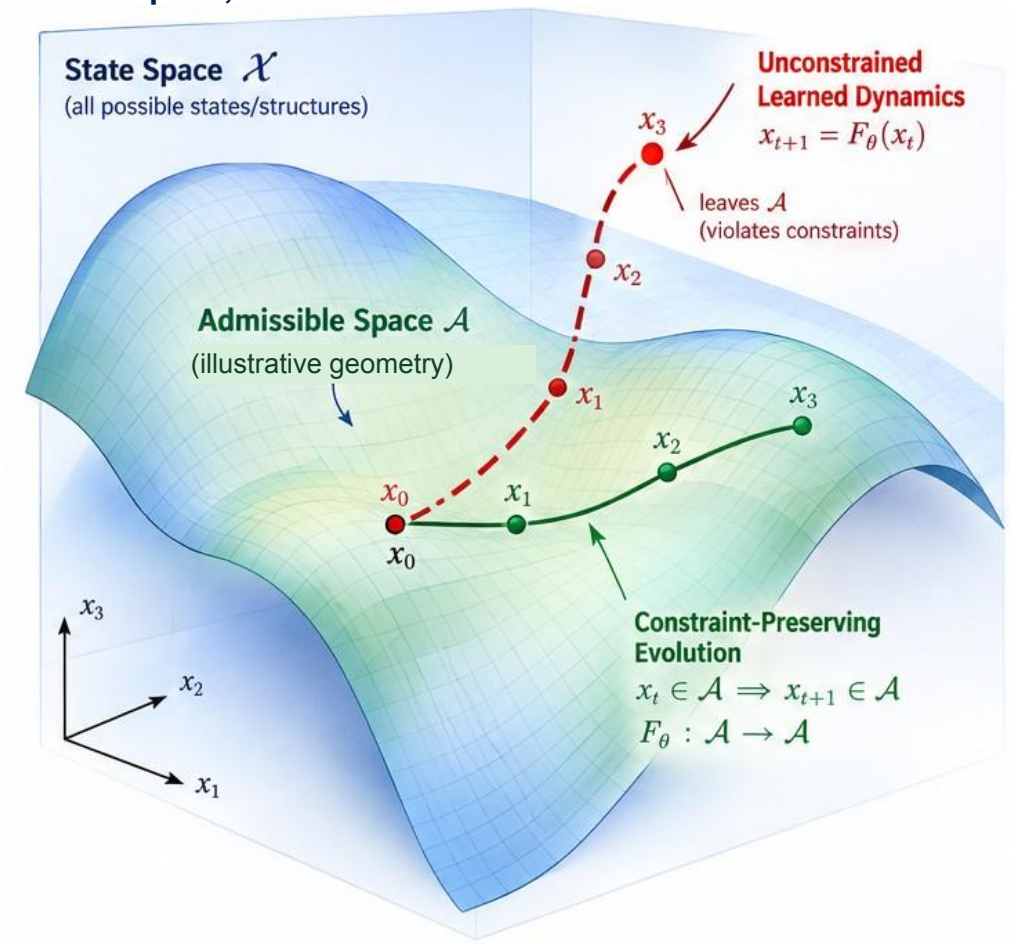

1. Admissible Space $\mathcal{A}$

The admissible space is defined by logical, physical, geometric, and engineering constraints.

$$\mathcal{A} = \left\{ x \in \mathcal{X} : C_i(x) = 0,\ G_j(x) \le 0,\ L_k(x) = \text{True} \right\}$$

- C_i : **equality constraints** (e.g., conservation laws, PDEs)
- G_j : **inequality constraints** (e.g., engineering limits, safety)
- L_k : **logical / discrete constraints** (e.g., topology, connectivity, type rules)

2. Constraint-Preserving Evolution

Fixed space and admissible set. Deterministic Fθ: $\mathcal{X} \to \mathcal{X}$.

$$\mathrm{F}_\theta(\mathcal{A}) \subseteq \mathcal{A}$$

$$\mathrm{x}_0 \in \mathcal{A} \ \Rightarrow\ \mathrm{F}_\theta{}^t(\mathrm{x}_0) \in \mathcal{A},\ \forall t \in \mathbb{N}_0$$

Proof by induction: admissible initial state + one-step closure.

3. Structure Generation under Constraints

Search for useful structures under specified constraints.

$$S^* = \arg\max_{S \in S_{\text{Admissible}}} U(S) \quad \text{or} \quad \text{Generate } S \text{ s.t. } S \models \Phi$$

(if the optimum is attained)

Intelligence = **Structure Generation under Constraints**

Examples of structures S

- Mathematical proof
- Program / algorithm
- Graph / knowledge structure
- Circuit / chip layout
- Physical field / simulation
- Control trajectory
- Molecule / material structure
- Multi-step plan / agent behavior

4. Domain-Specific Specification Φ

Conjoin applicable requirements under shared semantics.

$$\Phi = \Phi_{\text{logic}} \wedge \Phi_{\text{physics}} \wedge \Phi_{\text{geometry}} \wedge \Phi_{\text{engineering}}$$

- Φ_{logic} : formal logic, type system
- Φ_{physics} : physical laws, PDEs
- Φ_{geometry} : geometric/topological constraints
- $\Phi_{\text{engineering}}$: design rules, safety, manufacturability

EDA: preserve $\mathcal{A}_{\text{hard}}$ during search. Verify $\mathcal{A}_{\text{target}}$ before acceptance ($\mathcal{A}_{\text{target}} \subseteq \mathcal{A}_{\text{hard}}$).

Huawei Proprietary - Restricted Distribution

HUAWEI

Exact, Approximate and Semantic Preservation

Different guarantees require different contracts and evidence

Within one state space

Exact preservation

$$F(\mathcal{A}) \subseteq \mathcal{A}$$

$$\forall x \in \mathcal{A}: F(x) \in \mathcal{A}$$

Approximate preservation

$$\|r_i(F(x))\| \le \varepsilon_i, \quad x \in \mathcal{A}$$

One-step residual bounds require additional stability analysis for long-horizon guarantees.

Specify the norm, tolerance, applicable domain and time horizon.

Across representations

Semantic preservation

$$T: \mathcal{X}_s \to \mathcal{X}_t, \quad T(\mathcal{A}_s) \subseteq \mathcal{A}_t$$

$$x \models_s \varphi \;\Rightarrow\; T(x) \models_t \varphi^T$$

$$\forall x \in \mathcal{A}_s, \forall \varphi \in \Psi$$

Example: design facts in agent context

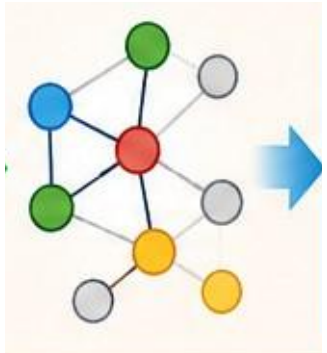

Structured context

U1 / NAND2
Connections: A, B → Y
Constraint: timing < 10 ps
Process: 5 nm
Source: spec_v1.2

LPVI Fig. 12 details and Fig. 9 context example. One-way property preservation also requires source-evidence checks.

HUAWEI

Formal Proof as a Reference Case

Legal transitions, completed proofs and kernel acceptance

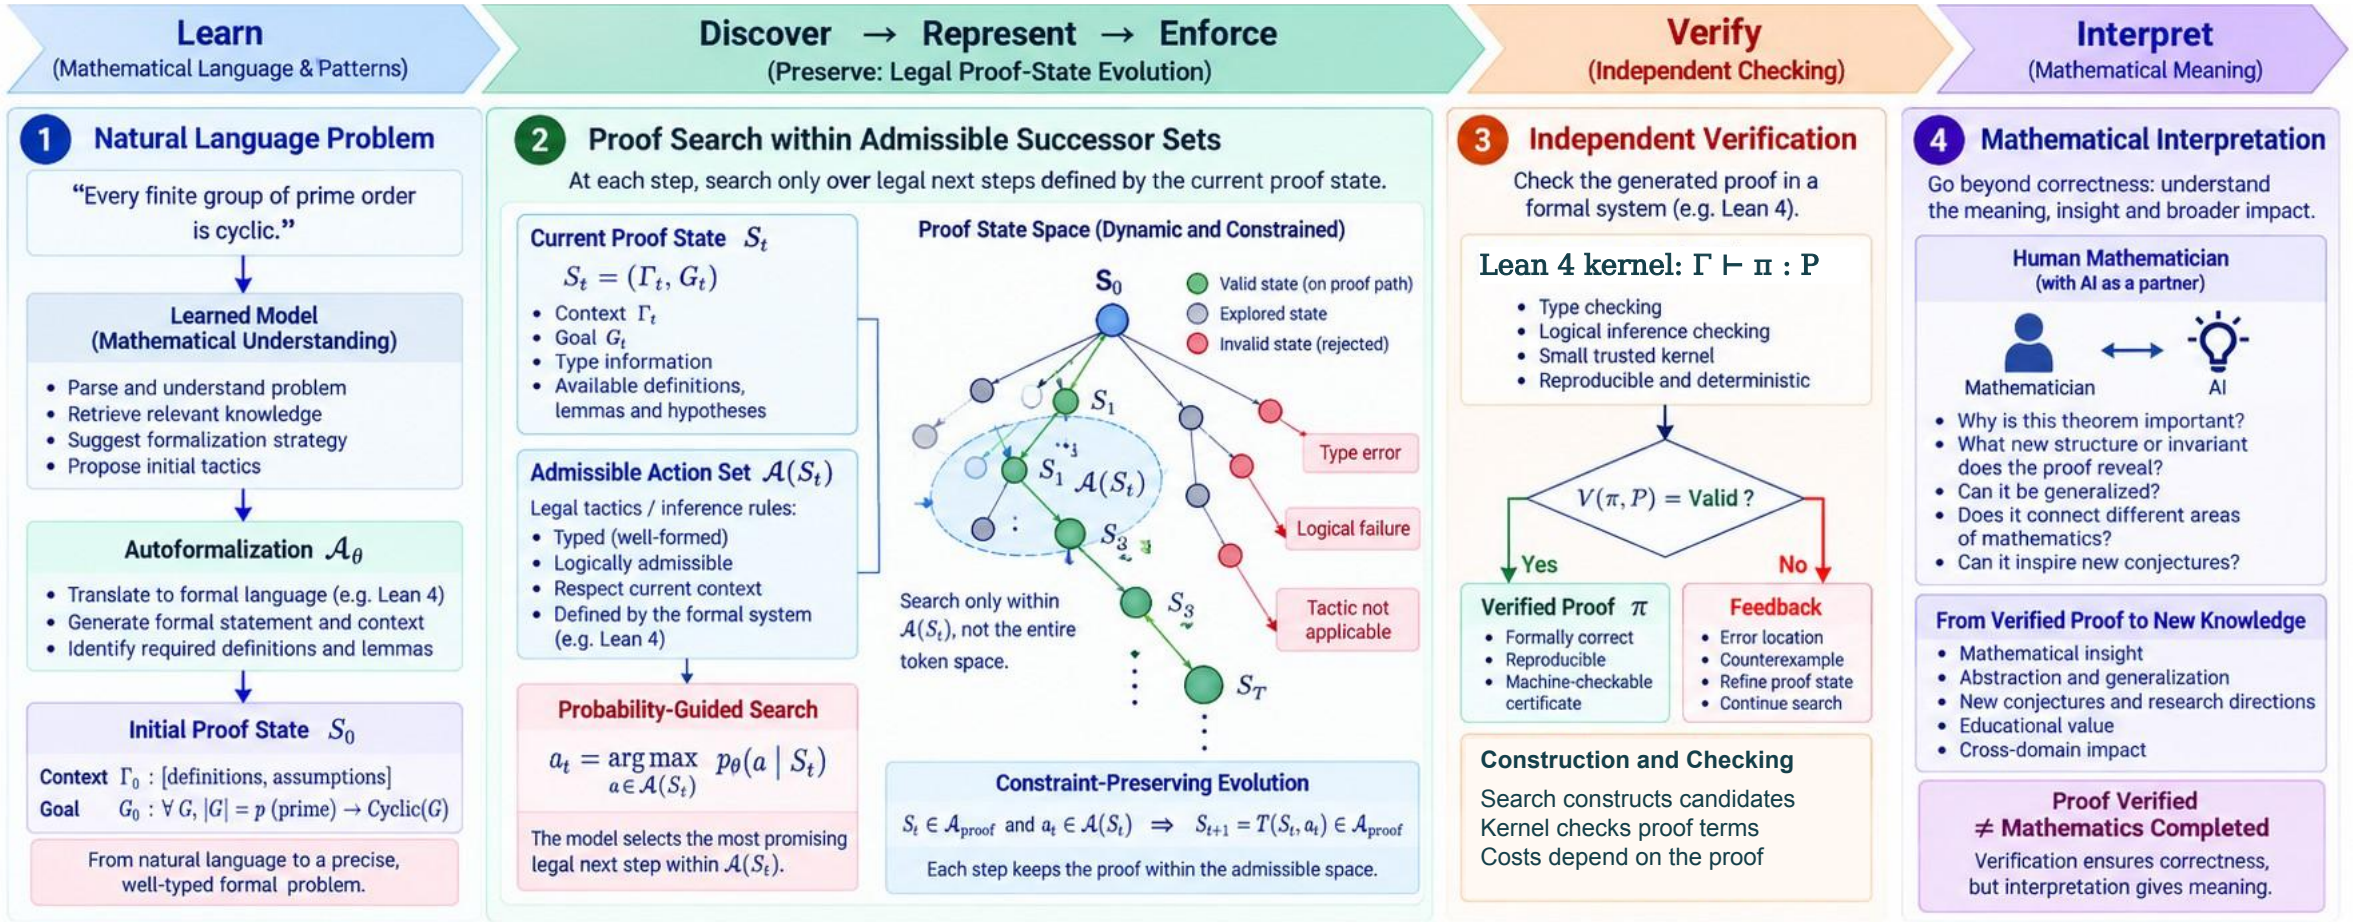

Legal steps do not guarantee proof completion.
Kernel acceptance does not establish faithful autoformalization.

HUAWEI

Semantic Contracts Across Representations

Target validity, specified property preservation and source evidence

LPVI Fig. 9 detail: data, ontology/KG, context and checks

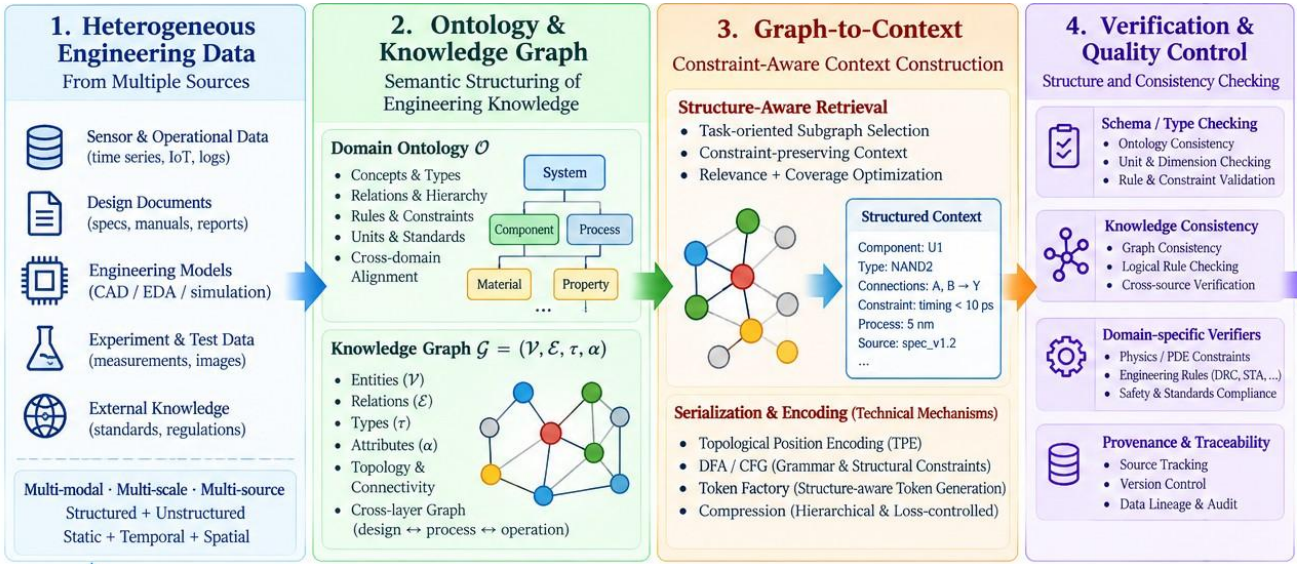

$$\mathrm{T} : \mathcal{X}_s \to \mathcal{X}_t$$

1. Target admissibility

$$\mathrm{T}(\mathcal{A}_s) \subseteq \mathcal{A}_t$$

2. Specified source properties

$$\forall x \in \mathcal{A}_s,\ \forall \varphi \in \Psi$$
$$x \models_s \varphi \Rightarrow \mathrm{T}(x) \models_t \varphi^{\mathrm{T}}$$

3. Source evidence

Check critical target assertions against supporting source records.

Illustrative failure: a requirement becomes a claimed result

Source requirement
Delay < 10 ps at SS, spec_v1.2

Generated assertion
“Timing passes at every corner.”

A schema can accept both records while missing the change in claim type, scope and evidential support.

Huawei Proprietary - Restricted Distribution

HUAWEI

Graph-to-Context under Evidence and Token Constraints

Context utility subject to critical constraints and evidence coverage

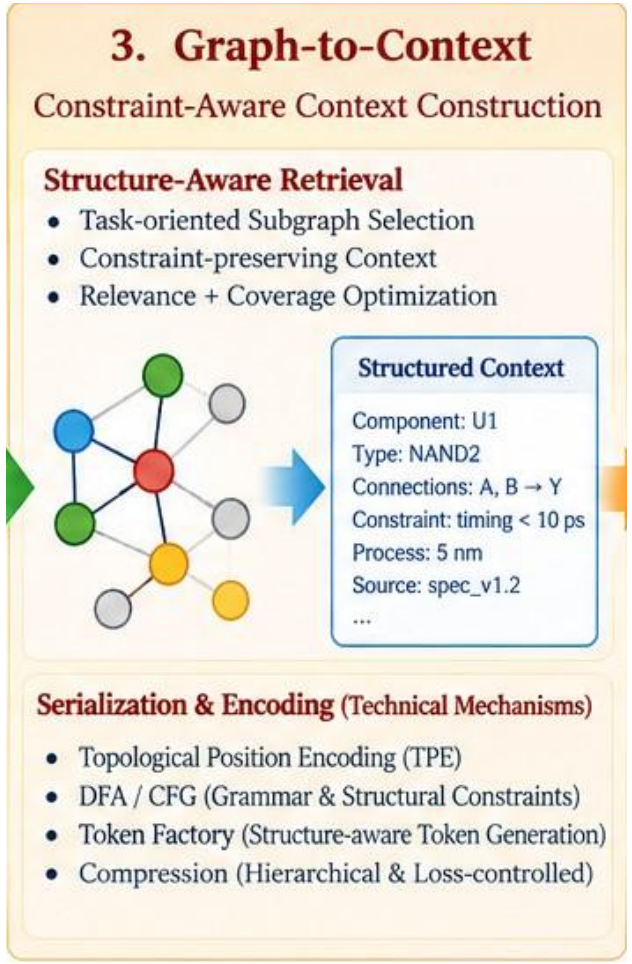

Task q: explain a timing result for a specified design snapshot and process corner

$$\max_C U(C; q)$$

$$\text{subject to } \text{Tokens}(C) \le B$$

$$C \models \Phi_{\text{critical}}$$

$$\text{Coverage}_{\Psi}(C, q) = 1$$

when complete critical-evidence coverage is required

Feasible context

Snapshot and specification
Corner, result and source IDs

Insufficient context

Expose missing evidence.
Report the budget conflict.

Relevance ranks candidates.
Evidence checks determine acceptability.

LPVI §9.3 and Fig. 9 detail. Coverage_{Ψ} follows the definition in the text.

Huawei Proprietary - Restricted Distribution

Token and Action Compilers

Shared specification, distinct generation and execution contracts

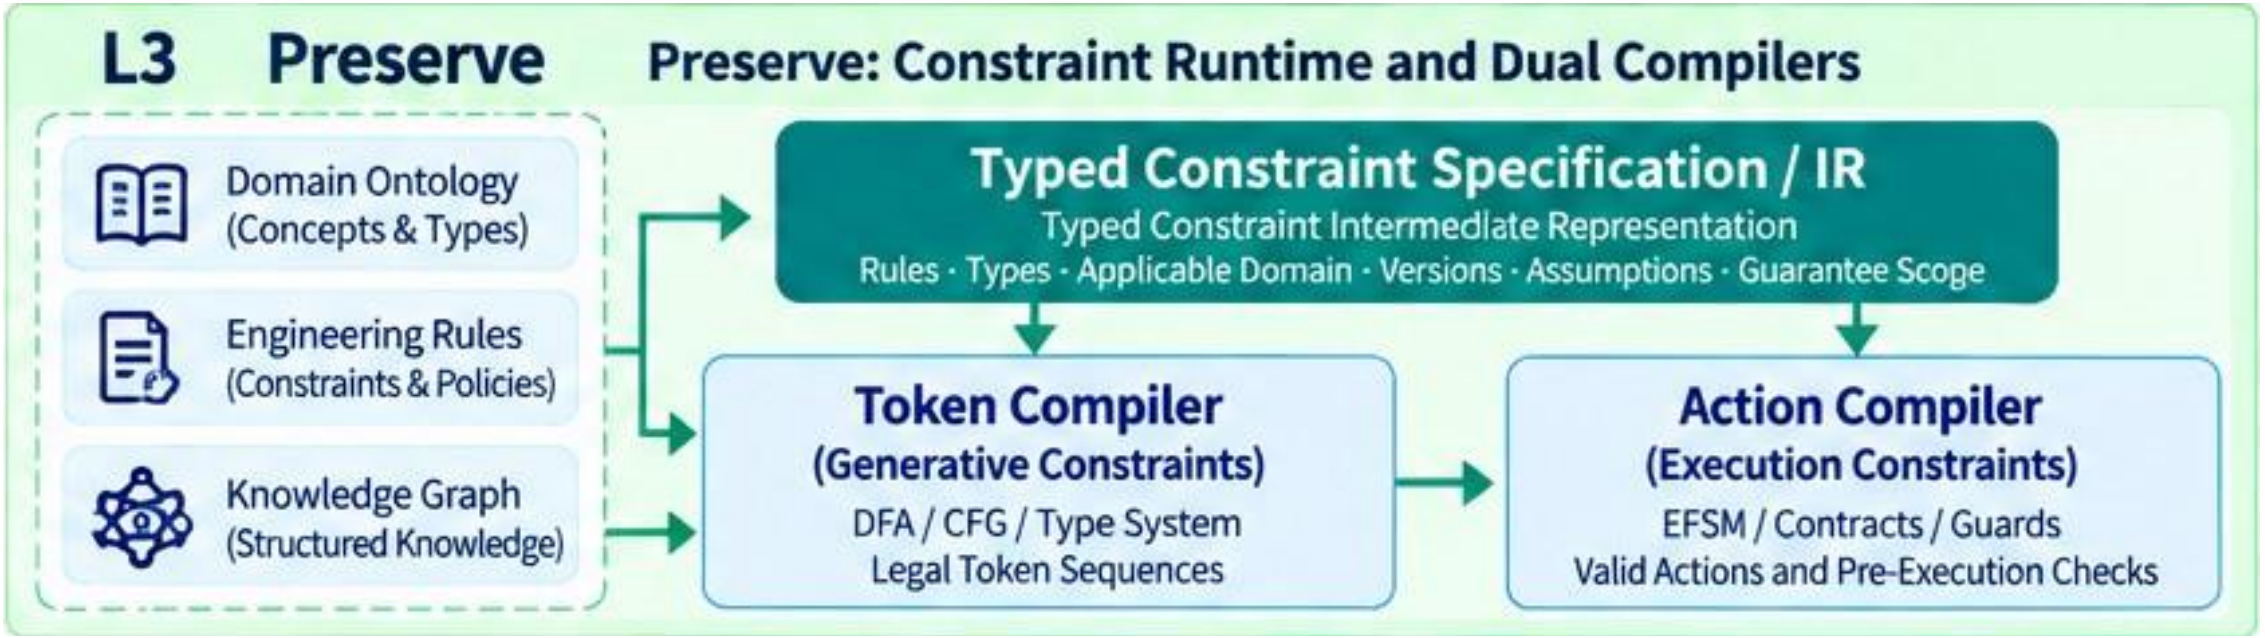

Token: language validity

$$\text{Keep } l_t(v) \text{ for } v \in V_{\text{legal}}(s_t).$$
$$\text{Mask other tokens with } -\infty.$$

Full token string. EOS only when accepting.

Action: execution permission

$$\text{Enabled}(q,s,a) = \text{ProtocolAllowed}(q,a) \wedge \text{Precondition}(s,a)$$

Guard pass does not establish postconditions.

HUAWEI

Agent Transitions and Compositional Verification

Policy–environment transitions and conditional local contracts



Policy and environment

$a_t \sim \pi_\theta(\cdot \mid S_t)$

$S_{t+1} \sim K(\cdot \mid S_t, a_t)$

$K^{\pi}(\mathcal{A} \mid S) = \int K(\mathcal{A} \mid S,a)\ \pi_\theta(da \mid S)$

Require $K^{\pi}(\mathcal{A} \mid S) = 1, \forall S \in \mathcal{A}$.

State includes the history needed by the policy.

Sequential partial correctness

$\{P\}\ A_1\ \{Q\},\ \ \{Q\}\ A_2\ \{R\}$

$\Rightarrow \{P\}\ A_1; A_2\ \{R\}$

Conditional on normal termination.
Concurrent writes can invalidate Q.

EDA: another agent replaces the netlist between two tool calls.

HUAWEI

Versioned Evidence and Transactional Acceptance

Evidence validity depends on the artifact, specification and environment



$$E = (H(S'), H(\Phi), V_{id}, \mathcal{E}, R, W), \quad R \in \{\text{Pass}, \text{Fail}, \text{Unknown}\}$$

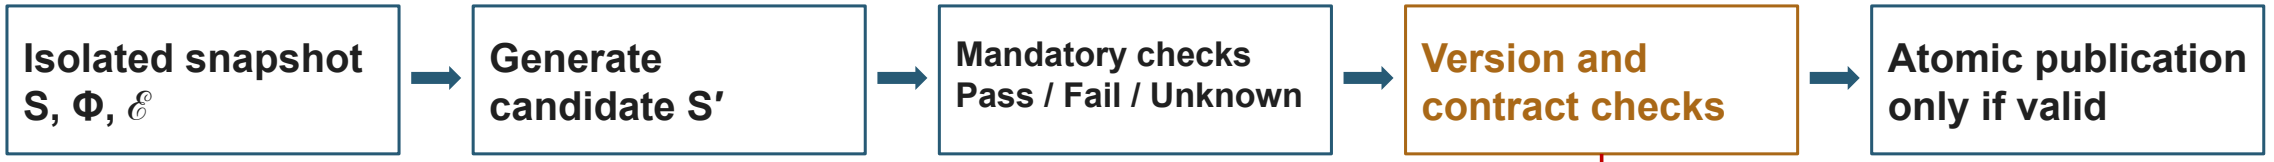

Illustrative change before publication

STA Pass under Φ_{v12} at corner SS.
Φ_{v13} changes a timing constraint.

Assess affected evidence.
Reverify the obligations
that no longer have support.

$$\text{Commit}(S') \Rightarrow \text{Verify}(S', \Phi; \mathcal{E}) = \text{Pass} \wedge \text{VersionValid} \wedge \text{ContractsCompatible}$$

HUAWEI

The Structural Intelligence Stack

Five deployment layers, four responsibilities and explicit trust boundaries

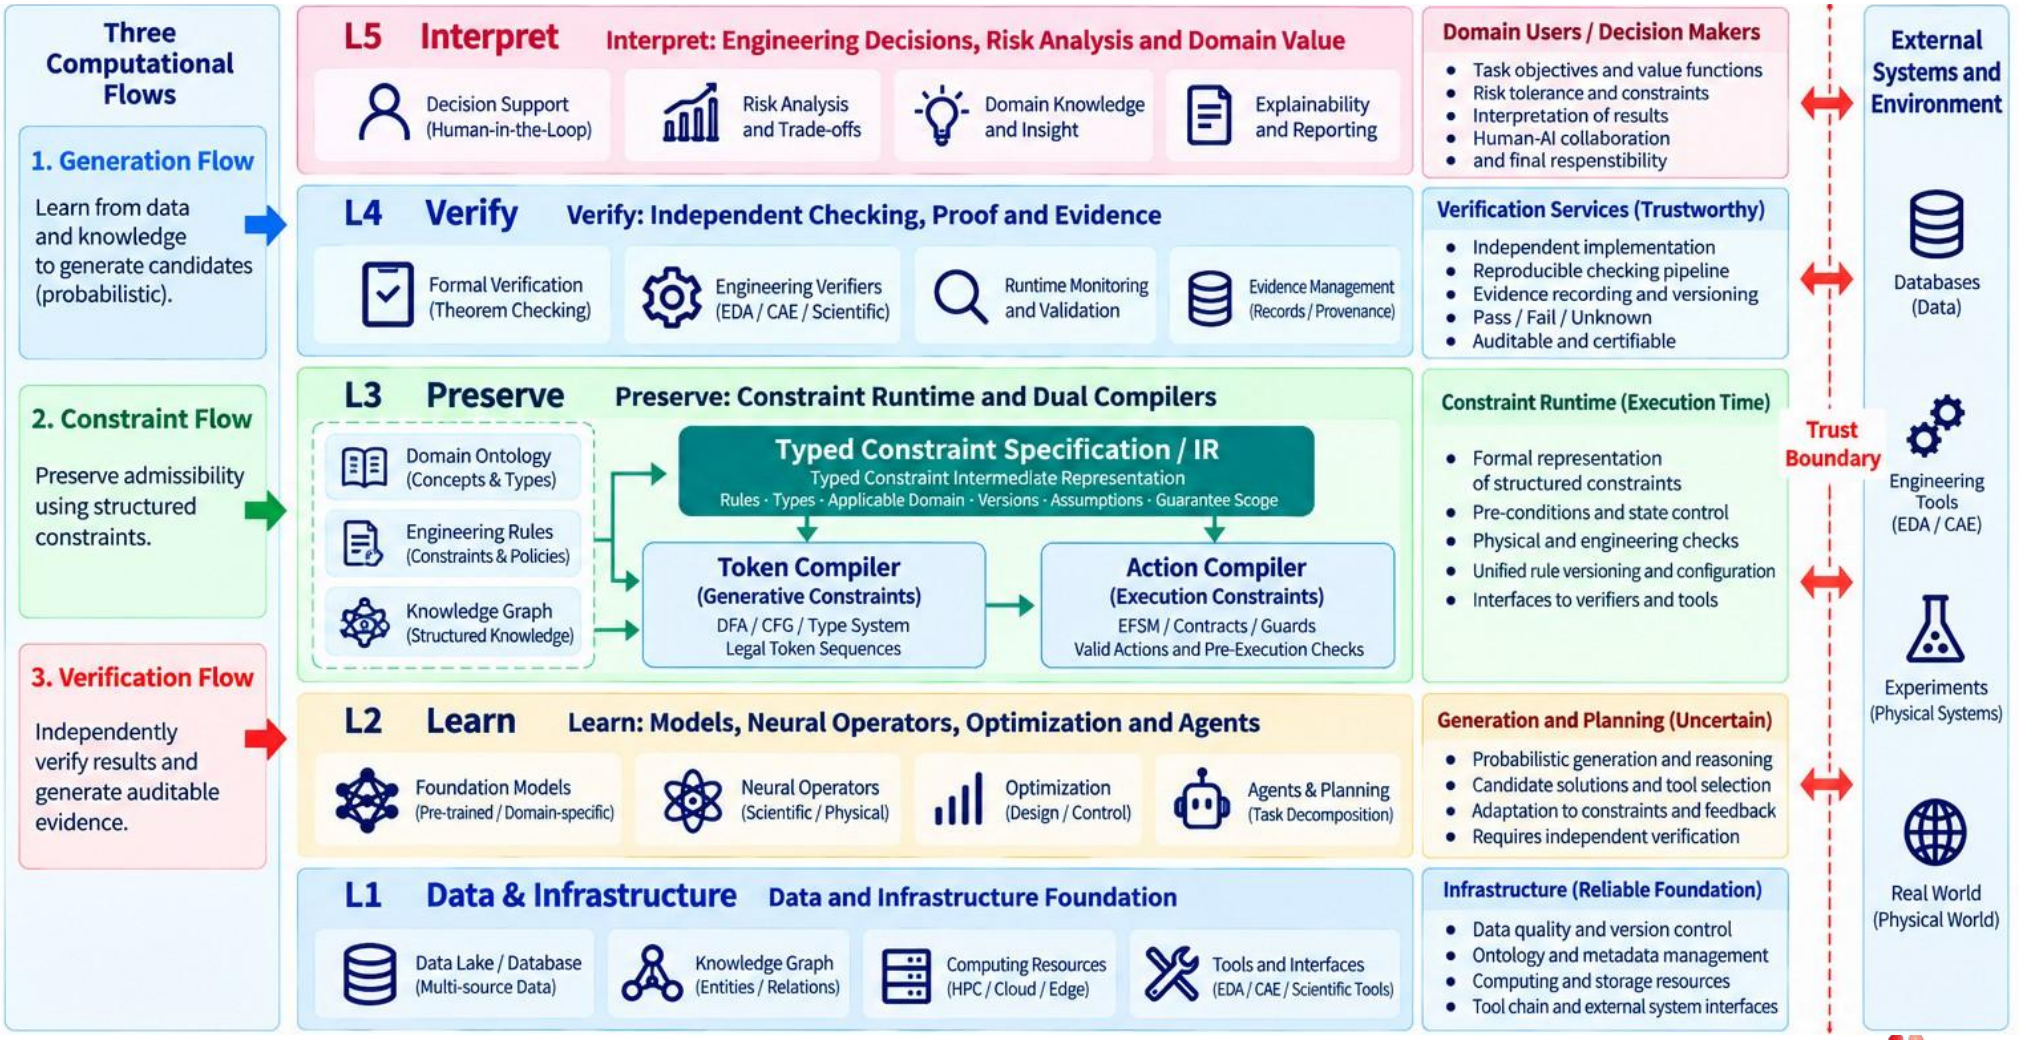

Huawei Proprietary - Restricted Distribution

Falsifiable Evaluation and the Cost of Trustworthiness

Proposed experiments: same EDA task set, data and generation budget



Configuration (cumulative)	Primary evaluation	Failure injection
Tuned retrieval-only baseline	Task utility and evidence fidelity	Missing corner or source
+ Soft constraints	Violation–quality trade-off	Conflicting requirements
+ Token / Action constraints	Language and protocol validity	Malformed call / forbidden order
+ Independent checks	Coverage and false acceptance	Timeout / skipped sign-off
+ Versioned evidence gate	Stale-evidence acceptance	Constraint update / concurrent write

Report false acceptance, coverage, Unknown outcomes and PPA/task utility.
Count latency, computation, human review and maintenance effort.

$$\min (C_{compute}, C_{verify}, C_{maintain}), \text{ subject to declared risk and latency targets}$$

Predeclare rejection criteria. Proposed design, no measured results.

HUAWEI

Open Problems in Structural Intelligence

Open research objectives, not a completed universal theory

R1 Structure Discovery
Discover unknown structures from data

Research Question:
How can we discover underlying geometric, topological, or symbolic structures from data, and integrate them with known priors?

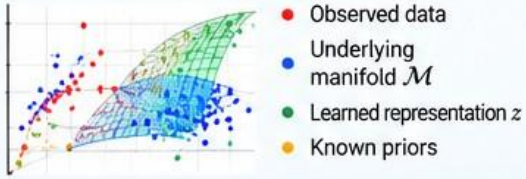

Hypothesis: $x \approx z \in \mathcal{M} \subset \mathcal{X}$

Methods: Manifold learning, spectral methods, TDA, neural-symbolic learning, physics-guided modeling.

Key Challenges: Data distribution shift, identifiability of structure, generalization and external knowledge integration.

R2 Constraint Preservation
Preserve key invariants in generation and evolution

Research Question:
How can we design generative and evolutionary operators that preserve key invariants in a specified admissible space?

Exact Preservation:
For a given admissible set $\mathcal{A}$ and operator F,

$$F(\mathcal{A}) \subseteq \mathcal{A}$$

Guarantee: applies to well-defined admissible spaces and operators.

Approximate Preservation:

$$\|r_i(F(x))\| \le e_i \quad (x \in \mathcal{A})$$

Note: Residual bounds do not automatically guarantee long-term stability; additional stability and error analysis are required.

Structure-Preserving Operators:
e.g., equivariant networks, symplectic integrators, divergence-free parameterizations, boundary embedding, continuous–discrete optimization.

R3 Compositional Verification
From local guarantees to system-level trust

Research Question:
How can we achieve scalable and reliable verification for systems composed of learned components, external tools, and complex environments?

Modular Verification:
$\{P\}\ A_1\ \{Q\},\ \{Q\}\ A_2\ \{R\}$
$\Rightarrow \{P\}\ A_1;\ A_2\ \{R\}$

Methods: Hoare Logic, assume–guarantee reasoning, model checking, SMT solving, proof certification.

Challenges: Component uncertainty, interface assumptions, environmental changes, scalability and computational cost.

Goal: Achieve compositional correctness under explicit assumptions, and extend to open environments and learning-enabled systems.

R4 Scientific Interpretation
From verifiable results to meaningful knowledge

Research Question:
How can we transform verified results into scientifically meaningful understanding and engineering insights?

Evidence to engineering value

Which verified trade-offs support the product objective?
What residual uncertainty remains?

Methods: Causal analysis, interpretable models, visualization, counterfactual reasoning, cross-modal mapping.

Challenges: Semantic definition, faithfulness of explanation, knowledge transferability, and human–AI collaboration.

Composability (Open Research Question)

How do exact, approximate and semantic guarantees compose?
Which evidence survives changes in design or specification?
How can evidence support decisions under a resource budget?

From plausible candidates to justified engineering acceptance

HUAWEI

Toward a Mathematical Semantics for LPVI

Open research sketch: categorical contracts and contextual evidence

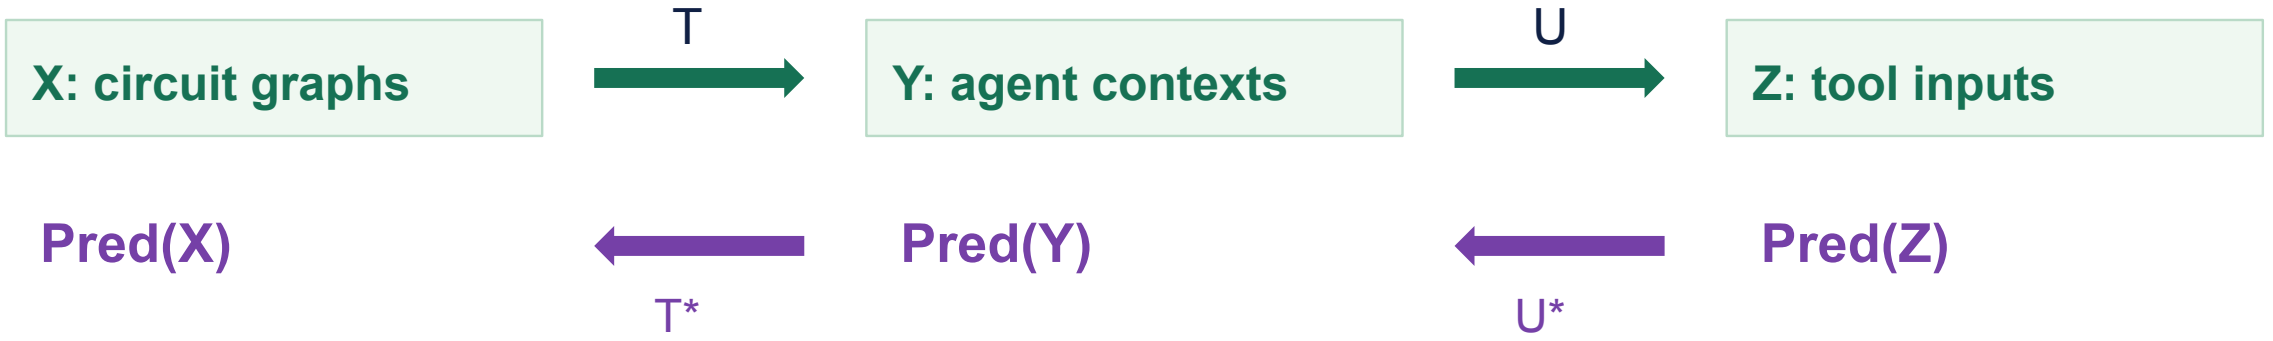

$$P \leq T^*Q, \quad Q \leq U^*R \quad \Longrightarrow \quad P \leq (U \circ T)^*R$$

Standard rule: ≤ means entailment; reindexing is monotone and functorial.

Indexed logic / fibrations

Which properties and proof evidence survive each representation change?

Sheaves / Grothendieck topoi

Can compatible local evidence glue? Define the site, coverage and restrictions.

Candidate semantics, not a completed LPVI theory. Probability and approximation need more.

HUAWEI

The Scientific Responsibility in the Age of AI

Knowledge access is becoming cheaper. Scientific judgment is not.

01

What is the right question?

Identify consequential scientific problems beyond benchmarks, fashions, and readily available answers.

02

What mathematical structure makes it tractable?

Discover invariants, constraints, abstractions, and representations that make a problem analyzable and computable.

03

What evidence establishes correctness?

Define proof obligations, independent checks, reproducible experiments, and the boundaries of each claim.

The deeper challenge is institutional, not only computational:

How should we organize research, sustain long horizons, validate results, and connect mathematical foundations to real systems?

From cheaper answers to better questions, meaningful structures, and reliable evidence.

Huawei Proprietary - Restricted Distribution

B1 Forward Invariance: What Closure Guarantees

Use after P4 / P5 · Mathematical backup for discussion

Fixed state space and admissible set

$$F : \mathcal{X} \to \mathcal{X}, \quad \mathcal{A} \subseteq \mathcal{X}, \quad x_0 \in \mathcal{A}$$

$$\forall x \in \mathcal{A} : \; F(x) \in \mathcal{A}$$

$$x_{t+1} = F(x_t) \quad \Longrightarrow \quad \forall t \geq 0 : \; x_t \in \mathcal{A}$$

Base: $x_0 \in \mathcal{A}$
Step: $x_t \in \mathcal{A}$ implies $F(x_t) \in \mathcal{A}$

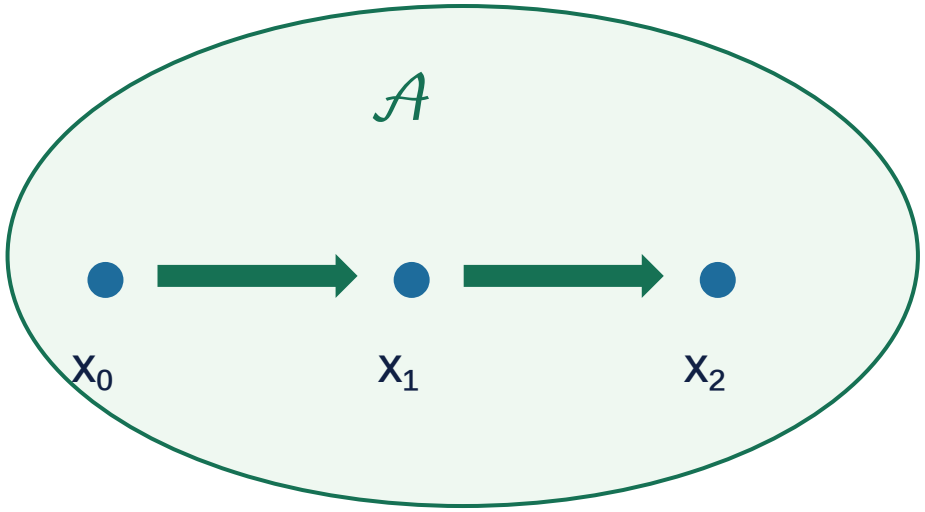

Nondeterministic updates

$$F(x) \subseteq \mathcal{A} \quad \text{for every } x \in \mathcal{A}$$

Closure proves membership at every step. Convergence and optimality need more.

Standard induction argument, shown explicitly. LPVI uses the obligation; it does not claim a new invariance theorem.

HUAWEI

B2 Quadratic Penalties and Exact Feasibility

Use after P2 / P4 · Mathematical backup for discussion

A scalar counterexample with an exact solution

$$\min_x (x-1)^2 \quad \text{subject to } x = 0$$

Finite quadratic penalty

$$J_\rho(x) = (x-1)^2 + \rho x^2, \quad \rho > 0$$

$$J'_\rho(x) = 2(x-1) + 2\rho x = 0$$

$$x_\rho = \frac{1}{1+\rho} \neq 0 \quad \text{for every finite } \rho$$

An exact-penalty exception

Replace ρx^2 with $\rho|x|$

At $x = 0$: $0 \in -2 + \rho[-1, 1]$

Thus $x^* = 0$ when $\rho \geq 2$
for this particular problem.

Small residuals and exact feasibility are different claims. State the tolerance.

Boyd & Vandenberghe, Convex Optimization, Ex. 5.14. Scalar derivation here; the threshold 2 is example-specific.

HUAWEI

B3 Semantic Preservation Across Representations

Use after P7 / P8 · Mathematical backup for discussion

$$T : \mathcal{X}_s \to \mathcal{X}_t, \quad T(\mathcal{A}_s) \subseteq \mathcal{A}_t$$

$$\tau : \Psi \to \text{target formulas}, \quad \varphi^T := \tau(\varphi)$$

Source x
SS corner, timing pass
spec_v1.2, known provenance



Target T(x)
Retain corner and version
“All corners pass” lacks support

$$\forall x \in \mathcal{A}_s,\ \forall \varphi \in \Psi : \quad x \models_s \varphi \implies T(x) \models_t \varphi^T$$

$$\text{Coverage}_\Psi(C) = 1 \quad \not\Longrightarrow \quad \text{all assertions in } C \text{ are true}$$

Selected-property preservation is one-way. Additional claims need evidence.

LPVI semantic interface proposal. Ψ, the property map τ, and both satisfaction relations must be specified.

HUAWEI

B4 Stochastic Preservation over a Finite Horizon

Use after P10 · Mathematical backup for discussion

$$K_t^{\pi}(B \mid s) = \int K_t(B \mid s, a)\, \pi_t(da \mid s)$$

Measurable kernels; state includes sufficient history for the Markov form

---

Exact almost-sure closure

$$\Pr(S_0 \in \mathcal{A}) = 1, \quad K_t^{\pi}(\mathcal{A} \mid s) = 1 \quad \forall s \in \mathcal{A},\ \forall t$$

$$\Longrightarrow \quad \Pr(S_0, \ldots, S_T \in \mathcal{A}) = 1$$

Conditional risk on surviving histories

$$\Pr(S_{t+1} \notin \mathcal{A} \mid H_t) \le \varepsilon_t \quad \text{on surviving histories}$$

$$\Pr(\exists t \le T : S_t \notin \mathcal{A}) \le \min\{1, \sum_{t=0}^{T-1} \varepsilon_t\}$$

Standard conditioning / first-exit argument. No independence needed. Finite T; sampled success does not prove closure.

HUAWEI

B5 Sequential Contracts and Their Composition

Use after P10 / P14 · Mathematical backup for discussion

Sequential composition with an explicit interface obligation

$$\frac{\{P\}\ A_1\ \{Q_1\} \quad Q_1 \Rightarrow P_2 \quad \{P_2\}\ A_2\ \{R\}}{\{P\}\ A_1; A_2\ \{R\}}$$

Example: placement succeeds, but routing still needs matching inputs

Q_1 must establish the required netlist, technology version and coordinate system. → P_2 refers to the exact artifact that the next tool actually consumes.

Partial correctness covers normal termination. Exceptions and interference need more.

Hoare (1969), sequence and consequence rules. Established logic, used here to expose cross-tool proof obligations.

HUAWEI

B6 Versioned Evidence and Atomic Acceptance

Use after P11 · Mathematical backup for discussion

$$E = (H(S), H(\Phi), V_{\mathrm{id}}, \mathcal{E}, R, W)$$

Snapshot S · specification Φ · verifier V · environment $\mathscr{E}$ · result R · witness/report W

Design / specification
Checker / environment
Relevant change: reassess validity

→

Evidence E
Pass / Fail / Unknown
Bound to versions

→

Atomic validation and commit

$$\mathrm{Commit} \implies \mathrm{AllMandatoryPass} \wedge \mathrm{VersionValid} \wedge \mathrm{ContractsCompatible}$$

These are necessary acceptance conditions. Hashes establish identity, not truth.

LPVI proposed schema. Relevant dependencies govern reuse. Atomic commit and serializable isolation are separate concerns.

HUAWEI

B7 Toward a Mathematical Semantics for LPVI

Open research sketch: categorical contracts and contextual evidence

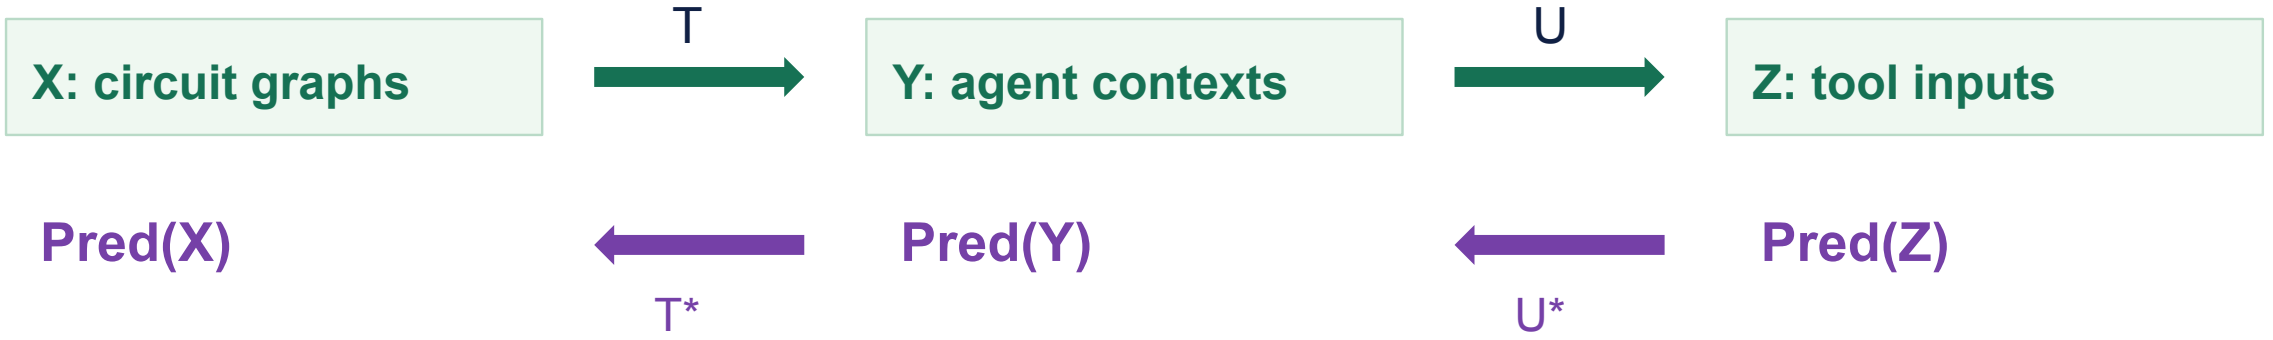

$$P \leq T^*Q, \quad Q \leq U^*R \quad \Longrightarrow \quad P \leq (U \circ T)^*R$$

Standard rule: ≤ means entailment; reindexing is monotone and functorial.

Indexed logic / fibrations

Which properties and proof evidence survive each representation change?

Sheaves / Grothendieck topoi

Can compatible local evidence glue? Define the site, coverage and restrictions.

Candidate semantics, not a completed LPVI theory. Probability and approximation need more.

HUAWEI