

The Grounding Gate: Admissibility and Replay Guarantees for AI-Driven Research

Thomas Dionysopoulos, CFA

Abstract

AI systems that generate computational pipelines from natural language may propose operations that are structurally valid but semantically unwarranted. Schema validation catches malformed proposals; it does not catch valid-but-wrong ones.

We present a *grounding gate*: a mandatory admissibility boundary between AI-proposed operations and deterministic execution. The system discovers which capabilities match the user’s terms by querying a live registry (236 capabilities), and a deterministic grounding function verifies that every name in the proposal has evidence in the discovery result. Names lacking evidence are rejected before execution, even if they name real capabilities. Admitted proposals execute deterministically, producing an 8-layer execution hash that decomposes provenance into semantic layers for fault localization without re-execution.

The capability registry separates description from identity: three layers (semantic, algebraic, implementation) determine a capability’s hash, while discovery metadata (aliases, tags) does not, allowing the registry to improve discoverability without invalidating prior execution hashes.

We evaluate on 30 prompts across 5 categories (4 strategy families, 9 metrics, 36 valid combinations). An unconstrained pipeline executes valid-but-unwarranted capabilities at **23.3%**; the grounded pipeline reduces this to **10.0%** (Fisher exact $p = 0.027$), eliminating them entirely on undiscoverable prompts (100%→0%). On adversarial prompts exploiting discovery-alias gaps, the grounded pipeline has a higher failure rate—a tradeoff inherent to constraining the admissible set. Repeated executions produce bit-identical hashes across 50 runs. Grounding overhead is under 14 ms.

1 Introduction

When a large language model proposes a computational pipeline, the resulting operations may not match what the user asked for. The operation names are real. The output is structurally valid. But the capabilities selected are not the ones the user’s request warranted.

Consider a researcher who prompts: “*Build a momentum strategy on equity futures, ranked by Sharpe ratio.*” A standard AI pipeline passes this to an LLM, which proposes `MOM_REV` (mean-reversion) as the strategy family and `SRP` as the ranking metric. Both are valid: `MOM_REV` is a registered strategy family; `SRP` computes the Sharpe ratio. Schema validation confirms the proposal is well-formed. Execution proceeds—and produces a mean-reversion strategy: the *opposite* computational signal from what the researcher requested. The two families produce computationally opposite outputs from the same infrastructure. *The output is correct in form and exactly wrong in substance.*

This failure is invisible to schema validation, constrained decoding, or any check that operates on the set of valid capabilities without reference to the user’s request. The proposed names are valid. The pipeline is well-typed. The output is structurally correct. The problem is that the user asked for momentum and received its inverse.

The root cause is architectural: there is no mandatory boundary between stochastic proposal and deterministic execution. The LLM’s output flows directly into the execution engine, constrained only by structural checks on its format (Figure 1, top row).

We present an architecture that introduces this boundary through a **grounding gate**. The system first discovers which capabilities match the user’s natural-language terms by querying a live registry (236 capabilities in our evaluation, spanning 4 strategy families and 9 evaluation metrics). The LLM then proposes a pipeline using discovered names. A deterministic grounding function verifies that every name in the proposal has evidence in the discovery result. If any name lacks evidence—even if it is a real capability—the proposal is rejected.

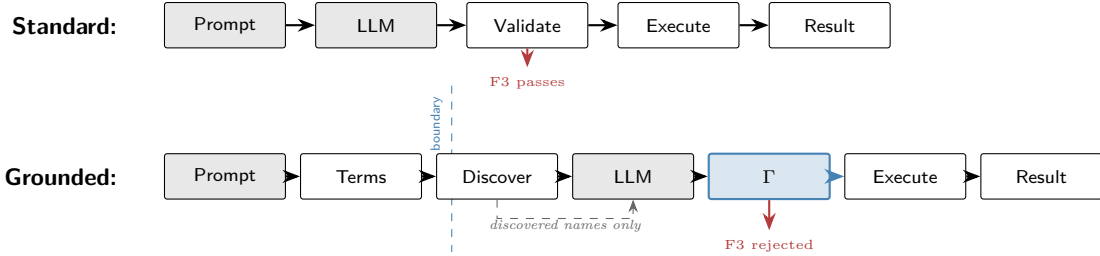


Figure 1: Pipeline comparison. Top: standard pipeline—the LLM proposes directly, validation passes F3 (valid-but-unwarranted). Bottom: grounded pipeline—discovery constrains the LLM’s vocabulary, grounding gate Γ rejects unwarranted proposals.

In the momentum example: if the LLM proposes `MOM_REV`, the grounding gate rejects it—`MOM_REV` is a real strategy family, but “momentum” discovers only `MOM_WZS`. Likewise, `CARRY` and `VOL_TGT` exist but are not discovered from this prompt.

This provides two properties. *Grounding safety*: no undiscovered capability name reaches execution. *Replay determinism*: identical grounded requests against identical data and registry produce bit-identical results, verifiable by an 8-layer hash that enables fault localization without re-execution.

The capability registry introduces a design principle we call *description/identity separation*. Each capability is hashed over semantic, algebraic, and implementation layers. A fourth layer—discovery metadata including aliases, descriptions, and tags—is excluded from the hash. This allows the registry to improve how capabilities are found without changing what they are, and without invalidating any prior execution hash.

We evaluate on 30 prompts in 5 categories, with 4 strategy families and 9 metrics (36 valid family \times metric pairs; discovery narrows to 1–3 per prompt). Against an unconstrained pipeline (schema validation only), valid-but-unwarranted executions occur at 23.3%; the grounded pipeline reduces this to 10.0% (Fisher exact $p = 0.027$), eliminating them entirely on undiscoverable prompts (100% \rightarrow 0%). On adversarial prompts exploiting discovery-alias gaps, the grounded pipeline performs worse—a tradeoff we discuss in §4.3. Repeated executions produce identical hashes. Grounding adds under 14 ms.

Contributions.

1. A **grounding gate** that rejects AI-proposed operations lacking evidence from a live capability registry, providing mandatory admissibility between stochastic reasoning and deterministic execution (§3.1).
2. **Description/identity separation** in capability hashing, where discovery metadata does not affect capability identity, allowing the registry to evolve its natural-language interface without invalidating prior execution hashes (§3.2).
3. **Fault-localizing execution hashes** that decompose end-to-end reproducibility into eight semantic sub-hashes, each covering a distinct computation layer (§3.3).

2 Background

We consider systems where an LLM generates computational pipelines from natural-language prompts. The pipeline specifies a computational family (e.g., momentum, carry, mean-reversion, or volatility targeting), a scoring metric, parameter ranges, and a data source. Our evaluation domain is systematic trading research; the families serve as deterministic, semantically distinct computational pipelines for testing admissibility—the paper makes no claims about financial performance. The system expands these into a grid of concrete configurations (morphisms), executes each, scores the results, and selects the best.

The design space separates two concerns: *what to run* (the LLM’s proposal) and *how to run it* (the deterministic execution engine). In existing systems, these concerns share a runtime with no formal boundary. An LLM may propose any name that appears in its training data, including names that are structurally valid but semantically unrelated to the user’s request.

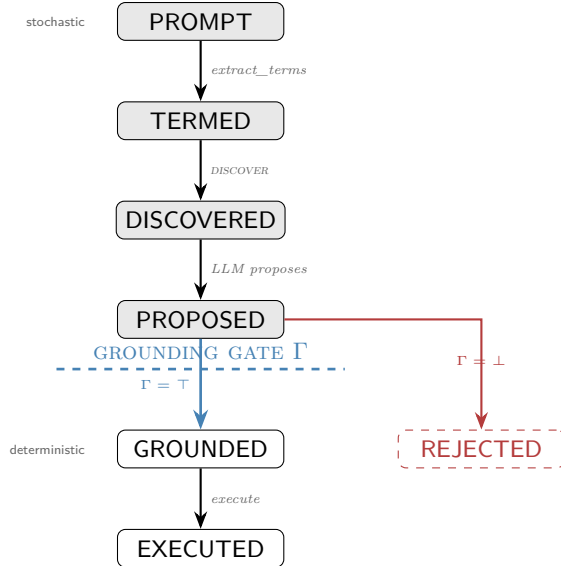



Figure 2: Pipeline state machine. The grounding gate Γ is the mandatory boundary between stochastic proposal and deterministic execution. Above the boundary, the LLM participates. Below, all functions are deterministic.



Adding aliases changes discovery, not identity

Figure 3: Capability hashing: three layers determine identity (h_{CPT}); the discovery layer is excluded. Changing aliases, descriptions, or tags does not change the capability hash.

We define this failure mode precisely:

Definition 1 (Valid-but-unwarranted execution). *An execution is valid-but-unwarranted if the proposed capability names pass schema validation (they exist in the system) but are not justified by the user’s natural-language request.*

This is distinct from hallucination (proposing nonexistent names) and from malformed output (unparseable structure). Standard validation catches both of those. It does not catch valid-but-unwarranted proposals because the validator has no model of what the user intended—only of what the system supports.

Notation and terminology. Table 1 summarizes the key terms, capability names, and hash layers used throughout the paper. Strategy families and metrics are the evaluation domain’s capability names; the architecture is domain-independent.

3 System Design

The system pipeline has three stages: *discover* capabilities from the user’s terms, *ground* the AI-proposed request against the discovery result, and *execute* the grounded request with full provenance hashing. We describe each stage and the properties it provides.

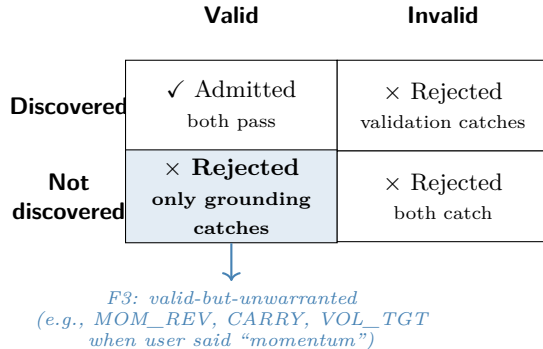


Figure 4: Grounding vs. validation. Schema validation partitions the family×metric space into valid (36 pairs) and invalid. The grounding gate further partitions valid pairs into discovered (1–3 per prompt) and not discovered. The bottom-left cell—valid but not discovered—is the valid-but-unwarranted region. Only grounding catches it.

3.1 Discovery and Grounding

Discovery. Given extracted terms from a natural-language prompt, AGENT-DISCOVER matches each term against the live capability registry. The registry contains 236 entries spanning operations (e.g., `dlog`, `wzs`), strategy families (`MOM_WZS`, `CARRY`, `MOM_REV`, `VOL_TGT`), and evaluation metrics (e.g., `SRP`, `MDD`). Matching uses a four-tier priority cascade:

1. *Exact*: the term equals a canonical name (confidence 1.0).
2. *Alias*: the term matches a registered alias (confidence 0.9). E.g., “sharpe” → `SRP`.
3. *Tag*: the term matches a semantic tag (confidence 0.7). E.g., “z-score” → `REL_MOM_Z`.
4. *Keyword*: the term appears as a substring in a name or description (confidence 0.5). E.g., “log returns” → `dlog`.

Terms with no match at any tier are placed in the UNRESOLVED set. The output is a discovery tuple containing a `MATCHES` list (each entry recording the term, the matched canonical name, the match tier, and the confidence) and an `OK` flag (true if at least one family and one metric were matched).

Discovery is deterministic: the same terms against the same registry produce the same result. The registry is queried live, not from a static allowlist—if a capability is added to the system, it becomes discoverable immediately without updating any configuration.

Grounding. The grounding function Γ takes a *request* (FPR—the AI-proposed pipeline specification, see Table 1) and the discovery result, and checks two conditions:

- The request’s family field must appear as the canonical name of a FAMILY-kind match in the discovery result.
- The request’s scoring metric field must appear as the canonical name of a METRIC-kind match in the discovery result.

If either condition fails, Γ returns \perp and execution cannot proceed. This is enforced architecturally: the runtime rejects any request that has not been grounded. A bare request without discovery evidence produces an error.

Property 1 (Grounding Safety). *For all requests q and discovery results r : if $q.family \notin names(r)$ or $q.metric \notin names(r)$, then $\Gamma(q, r) = \perp$.*

The contrapositive: if Γ admits a request, every capability name in it was discovered from the user’s terms.

Table 1: Notation summary. Families and metrics are evaluation-domain identifiers; the architecture is domain-independent.

Symbol	Meaning	Context
<i>Execution objects</i>		
Request (FPR)	AI-proposed pipeline specification	Input to grounding
Morphism	Concrete configuration with all parameters bound	Grid expansion
Capability entry	Registry entry with layered hash	CPT registry
Execution footprint	Record of a completed execution	Output provenance
<i>Evaluation-domain families (4)</i>		
MOM_WZS	Momentum (z-scored)	Opposite of MOM_REV
MOM_REV	Mean-reversion	Opposite of MOM_WZS
CARRY	Carry	Overlaps VOL_TGT in NL
VOL_TGT	Volatility targeting	Overlaps CARRY in NL
<i>Metrics (selected)</i>		
SRP	Sharpe ratio	9 metrics total
MDD	Maximum drawdown	
<i>Hash layers (8-layer execution hash)</i>		
Registry	Hash of all capability entries	Sensitive to op changes
Request	Hash of proposed family, metric, parameters	Sensitive to intent
Morphisms	Hashes of all grid-expanded candidates	Sensitive to grid
Score / Selection	Hash of scoring result / best candidate	Sensitive to output
Data	Hash of source data	Sensitive to input

Figure 2 shows the pipeline as a state machine. The critical boundary is between PROPOSED and GROUNDED. Above the boundary, the LLM participates—proposal is stochastic. Below the boundary, all functions are deterministic: morphism expansion, risk filtering, execution, scoring, and hashing depend only on the grounded request, the source data, and the registry state.

Why grounding is not validation. Schema validation checks whether a request’s field values belong to the set of real capabilities. Grounding checks whether those values have evidence from the user’s terms. In the motivating example, all four families—MOM_WZS, MOM_REV, CARRY, VOL_TGT—pass validation. But “momentum” discovers only MOM_WZS. Grounding rejects the other three, including MOM_REV: a real family built on the same infrastructure that produces the opposite computational signal. This distinction also separates grounding from constrained decoding, which restricts the LLM to the full set of 36 valid names but not to the per-prompt discovered subset (§5). Figure 4 illustrates this distinction.

3.2 Capability Identity

The capability registry assigns each entry a four-layer tuple. Three layers contribute to the capability’s identity hash:

- **Semantic** (:SEM): algebraic properties—pure, deterministic, commutative, associative, idempotent, orientation, stateful.
- **Algebraic** (:ALG): type signature—input types, output types, arity, reversibility.
- **Implementation** (:IMPL): execution details—engine, IR variant, fusability, execution path.

The identity hash is computed as SHA-256 over the canonical serialization of these three layers:

$$h_{\text{CPT}} = \text{SHA256}(\text{SEM} \parallel \text{ALG} \parallel \text{IMPL})$$

A fourth layer, **Discovery** (:DISC), contains aliases, descriptions, tags, usage examples, and deprecation status. This layer is intentionally excluded from the identity hash (Figure 3).

Table 2: Execution hash decomposition. Each sub-hash covers a distinct semantic layer, enabling fault localization without re-execution.

Layer	Sub-hash	Covers
Registry	DIC_HSH	Operation algebra (all capability hashes)
Request	FPR_HSH	Proposed family, metric, and parameters
Morphisms	MOR_HSHS	All grid-expanded candidate expressions
Plans	SRH_HSHS	Compiled execution plans
Artifacts	CMPN_HSHS	Output data artifacts
Score	SCR_HSH	Scoring result (metric values)
Selection	SEL_HSH	Best candidate identity
Data	DATA_HSH	Source data content

The exclusion is the **description/identity separation** principle: a capability’s identity is determined by what it *does* (semantic properties, algebraic structure, implementation), not by how it is *described* (natural-language metadata). Adding an alias like “log returns” \rightarrow `dlog` changes what the LLM can discover; it does not change what `dlog` computes.

This has a practical consequence. A researcher who executed a pipeline six months ago can verify their execution hash against the current registry. If only descriptions, aliases, or tags were updated—a common maintenance activity—the capability hashes are unchanged, the registry hash is unchanged, and the execution hash validates. We verify this empirically: adding aliases (e.g., “logarithmic returns” \rightarrow `dlog`), tags (e.g., `returns`, `log-space`), and extended descriptions to five operations used in the pipeline, rebuilding, and re-executing the same request produces identical hashes at every sub-hash layer (Table 5, DISC-layer row). Since `:DISC` is excluded from h_{CPT} , no change to discovery metadata propagates to h_C or to any downstream execution hash.

The full registry hash h_C is the SHA-256 of the sorted list of all capability hashes, producing a single fingerprint of the operation algebra. Every downstream hash includes h_C , making the entire provenance chain sensitive to changes in the capability set.

3.3 Execution and Provenance

Morphism expansion. A grounded request specifies a strategy family, scoring metric, and parameter ranges. The system expands parameter ranges into a grid via Cartesian product. Each grid point produces a *morphism*: a concrete configuration with all parameters bound (e.g., a specific signal window width, risk window, and leverage cap). Each morphism is canonicalized and hashed. The hash includes the canonical expression, bound parameters, and the registry hash h_C .

Risk filter. Each morphism is evaluated against four admissibility predicates: (1) all operations are resolved and plannable, (2) the source data file exists, (3) all parameters are non-negative, and (4) estimated leverage is within bounds. A morphism is *admissible* if all four predicates hold. Inadmissible morphisms are classified as rejected with the failing predicates listed.

Execution. Each admissible morphism is executed through a typed compilation pipeline (parse \rightarrow normalize \rightarrow canonicalize \rightarrow plan \rightarrow optimize \rightarrow execute). Execution produces an *execution footprint* recording the bound parameters, the computation graph (DAG of operations), and the output artifacts with content hashes.

Execution hash. The end-to-end execution hash decomposes into eight sub-hashes (Table 2):

The execution hash is the SHA-256 of all eight sub-hashes concatenated. This provides **replay determinism**: identical inputs and registry state produce identical hashes. When two hashes differ, comparing sub-hashes localizes the divergence. A change in the registry hash indicates the operation algebra was modified. A change in the data hash indicates the source data changed. A change in the score hash with everything else equal indicates the scoring computation changed. This fault-localization property allows diagnosis without re-execution.

Table 3: Grounding gate confusion matrix (96 tuples, 4 families, 6 discovery contexts). The gate is exact: it admits exactly the proposals whose names appear in the discovery result.

	Admitted	Rejected
Grounded (27)	27	0
Ungrounded (69)	0	69

Table 4: Replay determinism across 5 configurations \times 10 repetitions. Every sub-hash matches within configuration. Configs span all 4 families and 3 metrics.

Family / Metric	Parameters	Reps	Unique	Match
Momentum / Sharpe	window=25	10	1	✓
Carry / Sharpe	(defaults)	10	1	✓
Mean-reversion / Sharpe	window=25	10	1	✓
Vol-targeting / Volatility	window=60	10	1	✓
Momentum / Max-drawdown	window=50	10	1	✓

4 Evaluation

We evaluate four claims: (1) the gate is exact—zero false admits, zero false rejects on known grounding status, (2) execution hashes are deterministic across repeated runs and detect targeted changes, (3) an unconstrained pipeline silently executes valid-but-unwarranted capabilities that the grounded pipeline rejects, and (4) the grounding pipeline adds negligible overhead.

4.1 Grounding Gate Selectivity

Setup. We construct 96 request tuples across 6 discovery contexts spanning all 4 families: 27 with both family and metric present in the discovery result, and 69 with at least one field absent. The 69 ungrounded cases include: cross-family rejection (e.g., `CARRY` when the discovery terms are “momentum” and “sharpe”), semantic inversion (`MOM_REV` when only “momentum” is discovered), plausible but nonexistent names (`TREND_FOLLOW`, `TRACKING_ERR`), multi-family contexts where undiscovered families are rejected, and fabricated names.

Results. Table 3 reports the confusion matrix. The grounding gate admits all 27 grounded proposals and rejects all 69 ungrounded ones. **Zero false admits. Zero false rejects.**

The gate is exact: it admits exactly the proposals whose names appear in the discovery result. It is neither conservative (rejecting valid proposals) nor permissive (admitting ungrounded ones).

4.2 Replay Determinism and Drift Detection

Replay. We execute five request configurations 10 times each, recording the execution hash and all eight sub-hashes. Table 4 reports the results: all 50 runs produce identical top-level execution hashes within each configuration, implying identical sub-hash decompositions. Determinism holds across CSV parsing, multi-column frame operations, canonical serialization, and parameter grid expansion.

Drift detection. We introduce targeted changes and re-execute the same request (Table 5). All configurations use a 3-morphism sweep (signal window $\in \{25, 50, 100\}$) so that metric and data changes can affect selection. For data drift, we replace the file contents at a *stable path* so the request hash is unchanged—isolating the effect of data from the effect of the request.

4.3 Unconstrained Pipeline Comparison

Setup. We compare two pipelines on 30 prompts in five categories (6 straightforward, 8 confusable, 8 adversarial, 4 undiscoverable, 4 multi-family). Pipeline A is the grounded pipeline: discovery \rightarrow LLM

Table 5: Drift detection: which sub-hashes change under targeted modifications. “=” means unchanged; “ Δ ” means changed. The data-drift row uses a stable source path so the request hash is unchanged. The DISC-layer row validates description/identity separation: adding aliases, tags, and descriptions to five pipeline operations produces identical hashes at every layer.

Change type	Registry	Request	Morphisms	Plans	Artifacts	Score	Selection	Data
Scoring metric	=	Δ	=	=	=	Δ	Δ^*	=
Parameter value	=	Δ	Δ	Δ	Δ	Δ	Δ	=
Source data	=	=	=	=	Δ	Δ	Δ	Δ
DISC-layer metadata	=	=	=	=	=	=	=	=

*Selection hash changes when the best morphism differs under the new metric; with a single-morphism grid it would not.

Table 6: Failure taxonomy. F3–F5 are the failure modes that only the grounding gate catches; schema validation admits all three because the proposed names are real.

Code	Failure mode	Validation	Grounding	Risk
F0	Correct	—	—	—
F1	Malformed output	catches	catches	—
F2	Hallucinated name	catches	catches	—
F3	Valid-but-unwarranted	misses	catches	—
F4	Metric swap	misses	catches	—
F5	Partially grounded	misses	catches	—
F6	Parameter nonsense	partial	misses	catches

(sees discovered names only) \rightarrow grounding gate \rightarrow execution. Pipeline B is the unconstrained pipeline: the LLM receives all 36 valid family \times metric pairs, proposes a pipeline, validates against a JSON schema with enum-constrained fields, and executes directly.

The four strategy families—`MOM_WZS` (momentum), `MOM_REV` (mean-reversion), `CARRY` (carry), `VOL_TGT` (volatility targeting)—are semantically distinct but pairwise confusable. `MOM_WZS` and `MOM_REV` share identical infrastructure with opposite computational signals. `CARRY` and `VOL_TGT` share overlapping natural-language descriptions. Both pipelines use Claude Sonnet at temperature 0.7 with three repetitions per prompt.

Failure taxonomy. We classify each execution using a failure taxonomy (Table 6):

Results. Table 7 reports failure rates by category.

Pipeline A reduces valid-but-unwarranted executions from **23.3%** to **10.0%** overall (Fisher exact test, two-sided $p = 0.027$). The gate eliminates them entirely on undiscoverable prompts (0/12 vs. 12/12). Both pipelines achieve 0% on straightforward and multi-family prompts. The residual 10.0% in Pipeline A (9/90) arises from confusable and adversarial prompts where discovery returns multiple families and the LLM picks the wrong one—a within-discovery-set error that grounding alone cannot prevent.

The undiscoverable category is the cleanest contrast: prompts like “trend-following strategy” discover no family (“trend” is not an alias for any registered family). Pipeline A correctly rejects execution. Pipeline B picks a family from the full enum and executes silently (100% F3).

The adversarial tradeoff. On adversarial prompts, Pipeline A has a *higher* F3 rate (20.8%) than Pipeline B (12.5%). The strongest case is prompt **A02**: “momentum strategy—but bet against recent winners.” The user intends mean-reversion (`MOM_REV`), but “momentum” discovers only `MOM_WZS`. Pipeline A, constrained to the discovered set, proposes `MOM_WZS`—the opposite computational signal. Pipeline B, unconstrained, correctly proposes `MOM_REV` in all 3 repetitions. The grounding gate prevents undiscovered capabilities from reaching

Table 7: Valid-but-unwarranted (F3) failure rates by category and pipeline (90 trials each, Claude Sonnet, $T = 0.7$, 3 reps). F3 includes semantic inversions (e.g., proposing mean-reversion when momentum was intended). Classification is intent-based: the proposed capability does not match the human-annotated expected capability.

Category	Pipeline	N	F3	F3 rate
Straightforward (6)	A (grounded)	18	0	0%
	B (unconstrained)	18	0	0%
Confusable (8)	A (grounded)	24	4	16.7%
	B (unconstrained)	24	6	25.0%
Adversarial (8)	A (grounded)	24	5	20.8%
	B (unconstrained)	24	3	12.5%
Undiscoverable (4)	A (grounded)	12	0	0%
	B (unconstrained)	12	12	100%
Multi-family (4)	A (grounded)	12	0	0%
	B (unconstrained)	12	0	0%
Overall	A (grounded)	90	9	10.0%
	B (unconstrained)	90	21	23.3%

Table 8: Pipeline overhead (20 repetitions). Discovery and grounding add under 14 ms per request.

Stage	Time (ms)
AGENT-DISCOVER	13.6 \pm 0.9
AGENT-GROUND	13.2 \pm 0.6
AGENT-DRYRUN (incl. discovery + grounding)	13.4 \pm 0.8

execution, but it cannot prevent wrong selection when the correct capability falls outside the discovery system’s alias coverage.

Conversely, prompt **C02** (“momentum strategy that captures reversals”) shows the gate’s value on confusable prompts: Pipeline B proposes `MOM_REV` (the opposite signal) in all 3 repetitions, while Pipeline A, constrained to the discovered set (“momentum” \rightarrow `MOM_WZS` only), proposes the correct family. Schema validation cannot distinguish them—both are valid families, both produce well-formed output.

4.4 Overhead

Table 8 reports pipeline overhead measured over 20 repetitions. Discovery and grounding together add under 14 ms per request, dominated by process startup rather than computation.

The cost of the grounding gate is negligible relative to both LLM latency (typically 1–3 seconds) and full pipeline execution.

F3 residual. The grounded pipeline’s 10% uncaught error rate (Table 7) consists entirely of F3 (valid-but-unwarranted) failures—cases where the proposed capability name exists in the discovery set but does not match the user’s intent, including semantic inversions (e.g., proposing `MOM_WZS` when the user intended mean-reversion). The gate catches proposals referencing undiscovered capabilities but cannot distinguish correct from incorrect selection within the discovered set. Reducing F3 requires richer discovery semantics (e.g., intent-aware matching), not a stronger gate.

5 Related Work

Tool-augmented LLMs. ReAct [1], Toolformer [2], and frameworks like LangChain [3] enable LLMs to call external tools. The LLM directly selects which tool to invoke with no admission gate between selection

Table 9: Guarantee coverage across system categories. • = full, ◦ = partial, — = absent.

	<i>Admission gate</i>	<i>Evidence-based</i>	<i>Replay determ.</i>	<i>Fault localiz.</i>	<i>Desc./ident. sep.</i>	<i>Live registry</i>
Tool-augmented LLMs	—	—	—	—	—	◦
Structured output	◦	—	—	—	—	—
Experiment tracking	—	—	◦	◦	—	—
Content-addressed	—	—	•	◦	—	—
Typed APIs/DSLs	◦	—	—	—	—	—
This work	•	•	•	•	•	•

and execution. Gorilla [11] improves selection accuracy by retrieving API documentation, but the LLM selects freely from retrieved candidates—retrieval constrains context, not execution. A valid but wrong tool selection produces a silent failure in all these systems. Our grounding gate catches it.

Structured output and constrained decoding. OpenAI structured outputs [4] and Outlines [5] constrain LLM output to conform to a JSON schema or grammar. With enum fields, constrained decoding eliminates hallucinated names (F2). However, the enum encodes the full set of valid values—all 36 family×metric pairs in our evaluation—not the per-query discovered subset (typically 1–3 pairs). Constrained decoding prevents hallucinated names but not valid-but-unwarranted ones: the LLM may select `MOM_REV` for a momentum prompt because it is a valid enum value. The grounding gate restricts to the discovered subset (Table 6).

Experiment tracking. MLflow [6] and Weights & Biases [7] record execution artifacts and enable comparison across runs. They provide post-hoc observability: researchers can see what happened. They do not provide pre-execution admissibility: they cannot prevent an unwarranted execution from occurring. Our system provides both—the grounding gate prevents unwarranted execution, and the 8-layer execution hash provides post-hoc traceability.

Content-addressed computation. Nix [8] and Guix [9] content-address software builds, ensuring reproducibility through cryptographic hashing of all build inputs. Our execution hash applies the same principle to research pipelines. The key difference is in metadata treatment: Nix hashes *all* package metadata (descriptions, maintainer fields, etc.) into the derivation hash. Our description/identity separation explicitly excludes discovery metadata from the capability hash, allowing the natural-language interface to evolve independently of computation identity. We are not aware of content-addressed build systems that explicitly exclude a metadata layer from the derivation hash while retaining it for discovery.

Typed APIs and DSLs. Typed APIs and domain-specific languages constrain the space of valid programs. Languages like Stan [10] for probabilistic programming provide type safety and domain-specific validation. These systems prevent structurally invalid programs but do not address the prompt-to-program gap: they assume the programmer (or LLM) already knows the correct program to write. Our contribution is the discovery layer that constrains which programs the LLM *may* propose, based on evidence from the user’s natural-language terms.

Table 9 summarizes the guarantee coverage across system categories. To our knowledge, no prior system combines all six dimensions shown in Table 9 within a single execution architecture, though individual dimensions appear in existing work as noted above.

6 Discussion and Limitations

Grounding is syntactic, not semantic. The grounding gate checks name membership: is the proposed family name in the discovery result? It does not check whether the user’s research question is well-posed, whether the chosen family is the *best* match, or whether the parameter values make scientific sense. An LLM can still make bad research decisions within the grounded vocabulary. The gate prevents the LLM from selecting capabilities it was not offered; it does not prevent poor selection among offered capabilities.

Discovery coverage determines the gate’s ceiling. The grounding gate can only admit capabilities that the discovery system finds. When a user’s terms describe a capability indirectly—using infrastructure terminology rather than family names—discovery may fail to match, and the gate will block a correct proposal or force a wrong selection from the discovered set. Our evaluation confirms this: on adversarial prompts that exploit discovery-alias gaps, the grounded pipeline has a higher valid-but-unwarranted rate than the unconstrained pipeline (§4.3). Improving alias coverage in the discovery layer is the primary path to reducing this limitation, and it can be done without changing the grounding mechanism itself.

Single domain. We evaluate in systematic trading research. The architecture’s components—discovery, grounding, layered hashing—are domain-independent, but we have not demonstrated them in other domains. Generalizing to, e.g., molecular simulation or climate modeling pipelines would require building domain-specific capability registries. The grounding mechanism itself—match terms, check membership—is domain-independent.

Capability vocabulary. The evaluation registry contains 4 strategy families and 9 metrics (36 valid combinations). The architecture imposes no limit on capability count; the registry currently contains 236 entries spanning operations, families, metrics, signal blocks, and recipes. Discovery and grounding operate identically regardless of registry size. Evaluating with larger vocabularies—including vocabularies where discovery ambiguity increases—is future work.

Conservative registry fingerprint. The registry hash fingerprints the entire operation algebra. Adding an unrelated operation changes the registry hash, which cascades to the execution hash, even though the new operation has no effect on the pipeline. Per-operation dependency tracking (a finer-grained hash that includes only the operations actually used) is future work that would reduce false-positive drift alerts.

7 Conclusion

We presented a grounding gate that enforces a mandatory admissibility boundary between stochastic AI reasoning and deterministic execution. The gate rejects proposals whose capability names lack evidence from the user’s terms—including valid names that the system supports but the user did not request. On undiscoverable prompts, the gate eliminates valid-but-unwarranted executions entirely (100%→0%); on adversarial prompts that exploit discovery-alias gaps, it forces wrong selections—a tradeoff inherent to constraining the admissible set rather than the valid set.

Beyond admissibility, the system provides two complementary guarantees. Description/identity separation in the capability registry allows the discovery interface to evolve—adding aliases, tags, and descriptions—without invalidating prior execution hashes: we verify empirically that enriching five pipeline operations produces identical hashes at all layers. Eight-layer execution hashing decomposes end-to-end provenance into distinct semantic layers, providing deterministic replay across all 50 runs in our evaluation and enabling fault localization without re-execution.

The grounding gate addresses the gap between structural validity and warranted execution. Schema validation, constrained decoding, and typed APIs prevent malformed or nonexistent proposals; the grounding gate addresses the valid-but-unwarranted ones that those mechanisms miss. The overhead is under 14 ms.

A companion paper [12] formalizes the canonical execution semantics downstream of the grounding gate: the typed specification space, canonicalization, content-addressed hashing, and the 8-layer provenance chain that enables deterministic replay.

References

- [1] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao. ReAct: Synergizing reasoning and acting in language models. In *ICLR*, 2023.
- [2] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom. Toolformer: Language models can teach themselves to use tools. In *NeurIPS*, 2023.
- [3] H. Chase. LangChain [Software], 2023. <https://github.com/langchain-ai/langchain>.
- [4] OpenAI. Structured outputs. API documentation, 2024. <https://platform.openai.com/docs/guides/structured-outputs>.
- [5] B. T. Willard and R. Louf. Efficient guided generation for large language models. *arXiv:2307.09702*, 2023.
- [6] M. Zaharia, A. Chen, A. Davidson, A. Ghodsi, S. A. Hong, A. Konwinski, S. Murching, T. Nykodym, P. Ogilvie, M. Parkhe, F. Xie, and C. Zuber. Accelerating the machine learning lifecycle with MLflow. *IEEE Data Engineering Bulletin*, 41(4), 2018.
- [7] L. Biewald. Experiment tracking with Weights and Biases [Software], 2020. <https://wandb.ai>.
- [8] E. Dolstra, M. de Jonge, and E. Visser. Nix: A safe and policy-free system for software deployment. In *LISA*, 2004.
- [9] L. Courtès. Functional package management with Guix. In *European Lisp Symposium*, 2013.
- [10] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 2017.
- [11] S. G. Patil, T. Zhang, X. Wang, and J. E. Gonzalez. Gorilla: Large language model connected with massive APIs. *arXiv:2305.15334*, 2023.
- [12] T. Dionysopoulos. Canonical execution semantics for stochastic program generators. In preparation, 2026.

A Prompt Dataset

The comparison experiment (§4.3) uses 30 prompts organized into five categories:

- **Straightforward** (6): Clear term-to-capability mapping. E.g., “Momentum strategy on At.csv ranked by Sharpe.” (discovers MOM_WZS + SRP.)
- **Confusable** (8): Terms that overlap between families. E.g., “Momentum strategy that captures reversals.” (“momentum” discovers MOM_WZS; “reversals” is ambiguous between MOM_WZS and MOM_REV.)
- **Adversarial** (8): Language designed to elicit valid-but-wrong proposals. E.g., “Momentum strategy—but bet against recent winners.” (discovers MOM_WZS only; user intends MOM_REV.)
- **Undiscoverable** (4): Terms with no family match. E.g., “Trend-following strategy ranked by Sharpe.” (“trend” discovers no family.)
- **Multi-family** (4): Prompts explicitly naming two families. E.g., “Compare momentum and carry on At.csv.” (discovers MOM_WZS + CARRY.)

Each prompt is annotated with expected family, expected metric, discovery terms, and expected confusion type. The complete dataset is in `experiments/prompts_30.json`.

Table 10: Full failure taxonomy with detection mechanisms.

Code	Mode	Stage	Valid.	Ground.	Risk
F0	Correct	Completes	—	—	—
F1	Malformed	Parse	✓	✓	—
F2	Hallucinated name	Validation	✓	✓	—
F3	Valid-but-unwarranted	Grounding	×	✓	—
F4	Metric swap	Grounding	×	✓	—
F5	Partially grounded	Grounding	×	✓	—
F6	Parameter nonsense	Risk filter	○	×	✓
F7	Registry drift	Replay	×	×	×
F8	Provenance gap	Audit	×	×	×

B Failure Taxonomy Details

F3, F4, and F5 are the critical failure modes that distinguish grounding from validation. In all three cases, the proposed capability names are real and pass schema validation. The grounding gate catches proposals whose names lack evidence from the user’s discovery terms—eliminating F3 when the correct capability was never discovered (e.g., undiscoverable prompts: 100%→0%). When the correct and incorrect capabilities are both discovered (e.g., adversarial prompts where alias coverage is incomplete), the gate admits both; wrong selection within the discovered set is a residual failure the gate does not address (§4.3).

C Measured Overhead

Overhead measurements from 20 repetitions on the evaluation hardware:

- AGENT-DISCOVER: 13.6 ± 0.9 ms (median 13.0 ms, range 13–16 ms)
- AGENT-GROUND: 13.2 ± 0.6 ms (median 13.0 ms, range 13–15 ms)
- AGENT-DRYRUN (includes internal discovery + grounding): 13.4 ± 0.8 ms (median 13.0 ms, range 13–16 ms)

All three operations are dominated by process startup overhead (the BLISP subprocess). The computation itself is sub-millisecond. Discovery and grounding together add under 14 ms per request—negligible relative to LLM latency (typically 1–3 s).

D Artifact Repository

The companion code repository contains all experiment scripts, expected outputs, prompts, and verification tools. Within the repository:

- `artifacts/grounding-gate/` — experiment scripts and expected results
- `artifacts/shared/registry/` — registry snapshots (capability dictionary, alias tables, family definitions)
- `experiments/prompts_30.json` — the full prompt dataset

To verify all paper claims against existing results without re-running experiments:

```
bash reproducibility/verify_all.sh
```

Deterministic experiments (selectivity, replay, drift, overhead) require only the `blisp` binary and no API key. The comparison experiment (Table 7) requires an LLM API key and makes 180 API calls.