

Survival-Guided Discovery: Beyond Gradient Dependence

Imagination, Legality, and Continuation in the Grounded Inference Engine

Christopher P. Wendling • IntelliTrade Inc.

Abstract

Current artificial intelligence systems rely heavily on gradient-based optimization and explicit objective functions to guide search and learning. However, many real-world problems—particularly those involving construction, sparse rewards, or discontinuous solution spaces—do not provide usable gradients or smooth measures of progress. This paper introduces a general framework for search and discovery that does not depend on gradients. Instead, it is governed by legality, persistence, and continuation—principles already embedded in the Grounded Inference Engine (GIE). Within this framework, imagination is defined as a transform generator over structure space, constrained by legality and evaluated solely by survival. Gradients, when they exist, are treated as opportunistic accelerants rather than foundational requirements. **Survival-guided search is not an alternative to optimization—it is a strict generalization of it.** When gradients exist, they are used. When they do not, survival and continuation provide the governing signal. This formulation unifies evolutionary search, constraint satisfaction, and optimization under a single survival-based architecture.

1. Introduction

Many domains of interest—scientific discovery, physical construction, reasoning over abstract structures, and ARC-style intelligence tests—lack smooth objective functions. In such environments, partial progress may not be measurable, and correct solutions may only be identifiable at completion. This renders traditional optimization methods ineffective or inefficient, as they depend on continuous feedback signals.

The central question becomes: how can a system search effectively when no gradient exists?

This paper proposes that search need not be guided by proximity to a goal. Instead, it can be governed by structural viability: whether candidate constructions are legal, persist under constraints, and admit continued transformation. This approach is native to the GIE framework and extends its epistemic principles into action and construction spaces.

Survival-guided search is not an alternative to optimization—it is a strict generalization of it. When gradients exist, they are used. When they do not, survival and continuation provide the governing signal.

2. Imagination as a Computational Primitive

Within GIE, imagination is not treated as a creative or generative flourish, but as a formal operator:

Imagination is a transform generator over structure space, constrained by legality and evaluated solely by survival.

This reframes imagination as a proposal mechanism operating over a transform library. Each transform is constrained by legality—no contradictions, no invented structure, no violation of domain rules. Candidate structures are not trusted; they are tested.

Thus, imagination produces hypotheses in structural space, not answers.

3. The Failure of Gradient Dependence

Gradient-based systems assume three conditions: a known objective, a measurable distance to that objective, and a smooth path connecting initial and target states. These assumptions fail in many domains:

- Sparse reward environments provide no intermediate signal
- Combinatorial construction tasks lack continuity
- Many valid solutions are not reachable via incremental improvement

In such cases, optimization degenerates into random search or overfitting.

The failure is not computational—it is architectural. The system is attempting to optimize toward a target that cannot be locally perceived.

4. Survival as the Replacement for Distance

Instead of measuring distance to a goal, the proposed framework evaluates three properties:

- **Legality:** Does the structure obey constraints?
- **Persistence:** Does it remain stable under perturbation or exposure?
- **Continuation:** Does it admit further lawful transformation?

These properties define a scalar measure of viability independent of any explicit objective. Formally, a structure has *continuation* if it admits at least one non-zero-measure set of lawful transform sequences extending it without violating constraints or collapsing under evaluation. Simply: **a structure has continuation if it admits further legal evolution.**

Crucially, continuation introduces direction without requiring a goal. Structures that admit further development are favored over those that terminate. Search is therefore guided by the preservation of possibility rather than proximity to a target.

Concrete Example

Consider a block-stacking task where the goal height is unknown and no intermediate reward exists. Traditional optimization yields no gradient. Under survival-guided search, configurations are evaluated by stability and extendability. A single stable stack scores higher than scattered placements because it admits further lawful extension. Multi-block stacks are favored not by proximity to a goal, but by their capacity to continue. The goal state emerges as a terminal survivor.

5. Forward Survivability and Implicit Guidance

A stronger formulation evaluates not only current viability but expected future viability. For a candidate structure S , one can estimate the number of legal continuations and the survival rate of short transform sequences applied to S .

This produces a scalar measure of *forward survivability*. Structures with richer viable futures are preferred. Search is biased toward regions of structure space that support continued evolution.

Computational tractability.

Forward survivability need not be exhaustively computed. In practice, it is approximated via bounded sampling over short transform horizons (depth- k rollouts), local legality checks, and reuse of previously validated substructures. This yields a Monte Carlo estimate of continuation viability analogous to rollout-based evaluation in planning systems, but constrained to lawful transforms and survival criteria rather than reward accumulation.

The estimate is local, bounded, and survival-gated—not global or exhaustive.

This directly defuses the combinatorial explosion objection: the system does not enumerate all possible structures, it samples bounded continuations and discards those that fail survival.

6. Integration with GIE

This framework is a direct extension of GIE's epistemic principles:

- Voxel density corresponds to accumulated survival of structures
- Abstention corresponds to absence of viable continuation
- Assertion corresponds to survival under sufficient exposure
- The Entitlement Gate enforces legality and evidentiary sufficiency

Imagination becomes the proposal layer operating over this substrate. It generates candidate configurations in the Epistemic Survival Space (ESS), which are then filtered by survival criteria.

No proposal source earns trust. Only proposals earn survival.

7. Emergence and Use of Gradients

Gradients are not eliminated. They are reclassified.

Gradients do not create structure; structure creates gradients.

In many systems, gradients emerge only after sufficient structure has been discovered. Once a configuration begins to exhibit directional advantage—such as increased height in a stacking task or improved aerodynamic behavior in biological forms—a local slope becomes available. At that point, gradient-based optimization can accelerate progress.

This yields a two-regime system:

- **Discovery regime:** survival-guided, gradient-free
- **Exploitation regime:** gradient-assisted refinement

Importantly, gradients never override legality or survival constraints. And critically: **this framework does not discard optimization; it removes its necessity as a precondition for discovery.**

8. Correspondence with Natural Evolution

Natural evolution operates under the same principles:

- No predefined objective function
- No explicit measurement of distance to future traits
- Selection based on viability under current conditions

Structures arise through variation and are retained if they persist. Directionality emerges only when certain structures confer consistent advantage under environmental constraints.

Feathers, for example, did not originate as flight-optimized structures. They emerged as viable configurations—possibly for insulation or display—that later enabled aerodynamic advantages. Once such advantages appeared, evolutionary pressure amplified them. Gradients in natural systems are emergent, not foundational.

The proposed framework mirrors these constraints, not by imitation, but by adherence to the same underlying principles of viability and persistence.

9. Implications for Artificial Intelligence

This approach enables effective search in domains where traditional optimization fails: tasks with sparse or binary rewards, combinatorial construction problems, open-ended discovery environments, and non-stationary or partially observed systems.

It unifies three major search paradigms under a single survival-based architecture:

Evolutionary Search

Variation corresponds to transform proposals; selection corresponds to survival filtering under continuation. The survival filter is the selection operator.

Constraint Satisfaction

Legality enforces hard constraints at proposal time, eliminating invalid regions without requiring penalty functions. Infeasible regions are never entered.

Optimization

When a directional signal exists, it is incorporated as a bias over survivable structures, equivalent to local gradient ascent within the viable manifold.

Rather than optimizing answers, the system determines where answers are allowed to exist.

10. Conclusion

When gradients are absent, search must be guided by structure. When structure is absent, it must be discovered through survival.

Progress is not movement toward a goal. Progress is the preservation of viable transformation.

Within GIE, this yields a complete framework:

- Imagination generates candidate structures via lawful transforms
- Survival filters them against legality, persistence, and continuation
- Gradients, when present, accelerate refinement within the viable manifold

This formulation removes the dependency on predefined objectives and smooth reward landscapes, enabling systems to operate in environments where correctness cannot be incrementally measured.

It does not reject optimization. It generalizes it.

Only what survives continued exposure is real.

GIE Axiom: $\int_{\infty} : S \times \Omega \rightarrow S \cup \{\perp\}$ — S is structure $\rightarrow \int_{\infty}(S, \Omega) = S$