The Builder's View

A Primer on Constructive Accessibility and Building Inside Landscapes You Can Only Partially See

Michael Richard Nothem
Reachability Labs, Loveland, CO
mnothem@ReachabilityLab.org

March 2026

Abstract

For thirty years, the study of random constraint satisfaction has been conducted almost entirely from above. Researchers ask: when do solutions exist? How many are there? When do they fragment, freeze, or vanish? These are questions about the landscape as seen from an aerial survey. This primer introduces a different perspective: the view from the ground, where a construction crew pours a road one section at a time, can see only the terrain immediately ahead, and can never tear up concrete that has already set. From this perspective, a new kind of boundary emerges. Viable routes can be abundant—the terrain full of continuous paths from east to west—while the crew, constrained by its local instruments and irreversible commitments, can no longer reach any of them. The gap between what the aerial survey proves exists and what the crew can actually build is the existence–accessibility gap. The density at which this gap first opens is the constructive accessibility transition. This document explains, without equations, what the transition is, why it exists, what mechanism drives it, and why it was invisible until someone got down on the ground and started building.

1 Two ways to see terrain

Imagine a vast region of undeveloped land stretching between two distant points. A construction company has been contracted to build a continuous paved highway from the eastern edge to the western edge. Before construction begins, two teams study the terrain.

The first team is the **aerial survey crew**. They fly helicopters over the region, photograph the surface, measure elevations, characterize soil types, map drainage patterns. From all of this data, they can prove statistically that thousands of viable continuous routes exist from east to west. They know the terrain is navigable. But they have never traced any single route on the ground. Their knowledge is distributional—"this region has high route density, the connectivity is strong, the terrain is passable"—not constructive—"take this specific sequence of turns at these specific junctions." The survey is real science. It tells you what the terrain contains. It does not tell you how to build through it.

The second team is the **paving crew**. They are on the ground with heavy equipment. They pour concrete one section at a time, and once a section is poured, it is permanent. They cannot tear it up and re-route. They can see the terrain for roughly one section ahead—the slope, the soil stability, the drainage conditions—and they use their best instruments to pick the direction that looks most promising. Their instruments are good. They avoid steep grades, unstable soil, and obvious water crossings. They are skilled, methodical, and well-equipped.

The aerial survey and the paving crew are studying the same terrain. They are not seeing the same things.

The survey sees the distribution of viable routes. The crew discovers the distribution of *buildable* paths—routes that can actually be constructed by a crew that pours concrete forward, one section at a time, guided only by local instruments. These are different objects. A region can be full of viable routes (the survey says "navigable") while being impossible for the crew to build through without hitting an impassable formation (the crew says "inaccessible"). The discrepancy between what the survey proves exists and what the crew can actually construct is not a failure of either team. It is a real feature of the terrain that only becomes visible when someone actually tries to build through it.

This primer is about that discrepancy—what it is, when it appears, why it exists, and what it reveals about terrain that neither perspective alone can capture.

2 The classical view from above

Random constraint satisfaction is one of the most studied objects in mathematical science. The central question is deceptively simple: take a random formula—a collection of logical constraints on a set of variables—and ask whether there exists an assignment of values that satisfies every constraint simultaneously. This is the aerial survey: characterize what exists, without constructing anything.

The answer depends on how many constraints there are relative to the number of variables. This ratio is called the clause density, written α. At low density, constraints are sparse and satisfying assignments are abundant—the terrain is wide open, with routes everywhere. At high density, constraints are so numerous that no assignment can satisfy them all—the terrain is impassable. Between these extremes lies a sharp boundary—the satisfiability threshold—where the probability of a viable route existing drops from near one to near zero.

For random 3-SAT, the most studied model, this threshold is near $\alpha \approx 4.267$. Below this density, the aerial survey confirms that viable routes almost certainly exist. Above it, they almost certainly do not.

But the classical program discovered far more than a single boundary. As density increases toward the threshold, the structure of the viable routes undergoes a sequence of transformations. Routes that were once smoothly connected fragment into isolated corridors. Variables that could once take either value become frozen into a single assignment within each corridor. The geometry of the solution space transforms from a single connected highway network into an archipelago of narrow, rigid paths.

These discoveries are deep and beautiful. They are also, without exception, discoveries about the terrain as seen from the air. They describe what exists, how it is arranged, and when it changes character. They say nothing about what happens when a crew actually tries to build a road through the terrain, pouring concrete one section at a time.

3 The paving crew's predicament

The paving crew constructs a road variable by variable. At each step, the crew examines the current position—which variables have been set, which constraints are affected, what the local terrain looks like—and commits one more section of concrete in a chosen direction. The commitment is irreversible. The concrete sets. The crew cannot reconsider. They move forward until they either reach the western edge or reach a position where no viable direction remains.

The crew is not unskilled. They have a scoring instrument—a way of evaluating the immediate consequences of each possible next direction. A good instrument avoids directions that create immediate problems and prefers directions that keep future options open. In the specific process studied in the companion work (exact-all-local A0), the instrument counts post-move unit clauses: constraints that are one wrong step from being violated. The crew picks the direction that creates the fewest of these immediate hazards.

This instrument works. It manages local terrain effectively. The crew always has viable directions available. Hazard counts stay low throughout construction. By every local measure, the crew looks competent and the terrain looks cooperative.

And yet, above a certain terrain density, the crew fails approximately half the time.

The boundary where this happens is not at the point where viable routes vanish ($\alpha \approx 4.267$). It is not at the point where routes fragment ($\alpha \approx 3.86$). It is not at any boundary the aerial survey has identified. It is at $\alpha \approx 1.54$—deep in the regime where the aerial survey says the terrain is wide open, routes are astronomically abundant, and the landscape is thoroughly navigable.

This is the constructive accessibility transition. It is the terrain density at which a local crew with irreversible commitment loses the ability to reach destinations that still abundantly exist.

4 The existence–accessibility gap

The gap between what the survey proves exists and what the crew can build is not small. At $\alpha = 1.54$, the crew fails half the time. At $\alpha = 4.267$, viable routes vanish entirely. Between these two densities stretches a vast regime—covering most of the navigable landscape—where routes exist but are unreachable by this class of builders.

This gap can be measured precisely. At each density, one can plot both the probability that a viable route exists (the existence curve, from the aerial survey) and the probability that the crew successfully constructs one (the reachability curve). The horizontal separation between these two curves, measured at each viability level, is the existence–reachability gap profile $G(s)$.

At the 50% viability level, this gap is approximately 2.73—meaning the crew forfeits access to routes across a density range nearly three units wide.

The gap is not an artifact of a clumsy crew. The crew's instrument is optimal within its information class: it sees every legal direction, scores each one by the best available local criterion, and picks the best. A better instrument—one that probes the terrain more deeply, like ground-penetrating radar added to the surface instruments—can push the boundary higher, but cannot eliminate the gap entirely within the forward irreversible framework.

The gap is a feature of the relationship between the terrain and the crew. It measures the cost of building irreversibly with partial information in a landscape whose global structure the crew cannot see.

5 What the crew actually sees

The crew's world is small and local. At each section, they see the available directions, the immediate consequences of each (how many constraints become critical), and nothing else. They cannot see how many viable routes remain from their current position. They cannot see whether the terrain, given their current commitments, still contains any continuous path to the western edge. They cannot see whether the section they are about to pour will foreclose all remaining routes or leave thousands open.

This ignorance is not a design flaw. It is fundamental. The question "does a continuous viable route exist from this position to the western edge?" is itself a computationally hard problem—as hard as the original question about the full terrain. No efficient local instrument can answer it. The crew is not choosing to be ignorant. They are structurally unable to see the global property that determines their fate.

The companion empirical paper documented this invisibility quantitatively. Local diagnostic variables—the survivor fraction, the hazard count, the chosen-direction score—have essentially no predictive power for whether the crew will ultimately succeed or fail, until very late in construction. The crew's instruments read "clear terrain ahead" right up until the moment the road dead-ends.

This is not because the instruments are badly designed. It is because the property that controls success—whether a continuous route still exists from the current position to the destination—is a global property of the terrain-plus-committed-road pair, and no local instrument can detect it.

6 The mechanism: resolution boundaries

The analytical program that followed the empirical measurement discovered the specific mechanism by which the crew fails. It is not what anyone would have expected.

The crew does not fail because local hazards overwhelm them. Hazard counts stay low. Legal directions remain plentiful. The local terrain never deteriorates in any visible way.

The crew fails because of what happens inside its best-scoring directions.

At each junction, the crew identifies all directions with the best (lowest) hazard score. These directions form the minimum-score tier—the set of directions that look equally good by the crew's instruments. The crew picks uniformly at random from this tier. In a healthy junction, every direction in the tier leads to terrain where viable routes to the west still exist. In a dangerous junction, some directions in the tier lead to viable terrain and others lead to terrain where all routes have been cut off—but the instruments cannot tell which is which.

These dangerous junctions are not distinguishable from healthy junctions by any local measure. The directions that lead to viable terrain and the directions that lead to dead-end terrain have the same hazard score. They produce the same number of immediate constraint violations. They look identical through the crew's instruments. But they lead to fundamentally different global outcomes.

These are **resolution boundaries** of the scoring instrument: places where the instrument's ability to distinguish viable from non-viable paths breaks down. They are not defects in the terrain. They are not adversarially placed obstacles. They are natural features of the interaction between the terrain's global structure and the instrument's local resolution.

A resolution boundary exists wherever the viability of different successor directions diverges while their local scores remain tied. The divergence is real—one direction preserves viable routes, another eliminates them. The tie is also real—both directions create the same number of immediate hazards. The resolution boundary is the gap between these two realities.

Below $\alpha \approx 1.54$, resolution boundaries are rare. The crew navigates the terrain without encountering junctions where its best directions include hidden dead ends. Above $\alpha \approx 1.54$, resolution boundaries become frequent enough that the accumulated risk of choosing a dead-end direction becomes near-certain.

7 The walking dead: 99 blocks of perfect highway

The most striking feature of the constructive accessibility transition is not that the crew fails, but *when* the crew realizes it has failed.

Consider a crew that reaches a junction where three of its five best-scoring directions lead to viable terrain and two lead to dead ends. The instruments read the same for all five. The crew picks one of the two dead-end directions. The concrete sets.

What happens next is remarkable. The crew continues building. The terrain ahead looks fine. The instruments read "clear." The crew pours one section, then another, then another. Each section scores well. No hazards appear. The local view is healthy. The crew builds 99 more sections of perfect, flat highway.

Then, on section 100, the road abruptly dead-ends at an impassable geological formation. There is no way forward. The crew jams.

But the fatal mistake was not made at section 100. The fatal mistake was made 99 sections ago, at the junction where the crew chose the dead-end direction. For 99 sections, the crew was building a road to nowhere. Their instruments said everything was fine, but the committed road had already lost access to every viable route to the western edge. The crew was *walking* *dead*—locally legal, globally doomed—for 99 sections without knowing it.

This is not a rare pathology. In the measured data at system size 1000 near the constructive boundary, the median walking-dead interval exceeds 50 construction steps. The crew builds at least 5% of the road after the fatal mistake and before the visible failure. The process looks healthy for a substantial stretch after it is already irrecoverable.

The walking-dead interval is the temporal signature of the resolution boundary mechanism. Because the fatal choice and the visible failure are separated in time, local diagnostics cannot detect the failure when it happens. They detect it only when the consequences finally become locally visible—long after the irreversible commitment has been made.

8 Why the crew fails: it is the distribution, not the total

The resolution boundary mechanism raises an immediate question: what determines whether a junction is dangerous? The answer turns out to be surprisingly specific, and it has nothing to do with how many viable routes remain overall.

At each junction, the crew's viable routes are distributed across the available best-scoring directions. In a healthy junction, the routes are spread broadly: every direction that looks good by the instruments actually leads to terrain where viable routes continue. Pick any of them and the road stays alive.

In a dangerous junction, the routes have concentrated. A few of the best-scoring directions carry all of the remaining viable paths. The others carry none. But the instruments read the same for every direction in the tier, because the instruments measure local hazard, not the global distribution of viable futures behind each choice.

This is the critical distinction: the total number of viable routes from a position can still be enormous while the *distribution* of those routes across same-score directions has collapsed. The crew is not running out of routes. It is running out of directions that lead to routes—and it cannot tell the difference.

Below the constructive boundary, routes are distributed broadly enough that every best-scoring direction is safe. Above it, concentration has progressed to the point where some best-scoring directions are empty, and the crew's probability of stumbling into one accumulates across junctions until failure is near-certain.

9 The better instrument

Return to the paving crew. Their standard instruments measure surface properties: slope, soil stability, drainage. These are real measurements of real terrain features. But they are shallow—they see one section ahead and no deeper.

Now imagine upgrading the crew with ground-penetrating radar. The radar can detect subsurface formations—buried rock layers, hidden fault lines, underground voids—that the surface instruments miss. With radar, some junctions that were previously ambiguous become resolvable. A direction that looked identical to the others by surface instruments alone is revealed by the radar to have a subsurface formation blocking viable routes two miles ahead. The crew avoids it.

The companion empirical work demonstrates exactly this effect. A stronger instrument—one that propagates constraint information deeper than the bare surface score—eliminates the entire failure window. The crew builds successfully through terrain densities where the basic crew fails half the time.

The radar does not change the terrain. The subsurface formations are still there. What changes is the crew's ability to see them. The boundary moves because the crew's resolution improved, not because the landscape changed.

This is why the boundary is not a fixed property of the terrain. A different instrument with different depth of resolution would encounter the boundary at a different density. An instrument with perfect subsurface resolution—one that could detect every hidden formation before committing—would encounter no boundary at all below the point where routes truly vanish. The boundary lives in the relationship between the terrain's hidden structure and the instrument's ability to resolve it.

10 A new kind of boundary

The classical boundaries in random constraint satisfaction are properties of the terrain alone. The satisfiability threshold is where viable routes vanish. The clustering threshold is where routes fragment. The freezing threshold is where variables become rigid within each corridor. These boundaries exist independently of any crew or instrument. They are features of the territory.

The constructive accessibility transition is different. It is a property of the crew-terrain interaction. It depends on the terrain, the crew's instrument, and the gap between what the instrument can see and what determines the crew's fate. Change the instrument and the boundary moves. Change the terrain and the boundary moves. The transition exists in the relationship, not in either component alone.

This means the constructive accessibility transition is not a new fact about random 3-SAT. It is a new fact about what it means to *build inside* a structured landscape under irreversible commitment with partial information.

The distinction matters because it opens a different kind of question. The classical program asks: what is the terrain? The constructive accessibility program asks: what can a situated crew actually do inside the terrain? These are complementary questions. Both have precise formulations. Both exhibit sharp threshold phenomena. Both reveal real structure. But they reveal *different* structure, and neither subsumes the other.

The aerial survey sees features the crew cannot: the global distribution of routes, the clustering geometry, the freezing of variables across the landscape. The crew sees features the survey cannot: the resolution boundaries, the concentration of viable futures across locally identical choices, the long stretches of doomed highway separating hidden failure from visible jam.

A complete science of structured landscapes requires both views.

11 What this means beyond SAT

Random 3-SAT is the first calibration site for this kind of measurement, but the objects discovered there are not specific to SAT.

Any system where a process builds forward under irreversible commitment, guided by local instruments, inside a landscape with hidden global structure, is a candidate for the same phenomenon. The ingredients are general: a builder that commits, an instrument that sees locally, a landscape whose global structure is finer-grained than the instrument's resolution, and a gap between what the instrument reports and what determines whether the builder can still reach a valid outcome.

Compilers allocate registers under accumulating constraints. Schedulers commit resources one step at a time. Planners lock in actions that foreclose alternatives. Routing algorithms pour paths through congested networks. In each case, the process moves forward, the commitments are expensive or impossible to reverse, local metrics may look healthy, and the question of whether a valid completion still exists from the current position is a global property that no local instrument can efficiently answer.

The constructive accessibility transition says: there is a regime where the landscape is still full of valid outcomes, but the process has already lost the ability to reach any of them. The walking-dead interval says: the process will not know this has happened until much later. And the instrument-dependence of the boundary says: stronger instruments can push the boundary further, but cannot eliminate it entirely within the forward irreversible framework.

These are not analogies. They are structural claims about any builder-terrain pair with the right ingredients. Random 3-SAT provides the first clean measurement because it offers something most applied domains do not: a perfect oracle that can verify, for any partial construction, whether viable completions still exist. That oracle is what makes the walking-dead interval measurable, the resolution boundaries detectable, and the gap between existence and reachability precisely quantifiable.

The measurement is done in SAT. The object it measures is general.