PERSPECTIVE

Why Construction AI Must Understand Where

AI cannot reliably reason about project impact, readiness, or risk if location remains inconsistent text scattered across disconnected systems.

By Mark Klusza · CTO, Inertia Systems, Inc. · September 6, 2026 · 5 min read



Construction AI is often asked to summarize documents, extract information, classify records, and recommend actions. Those capabilities are useful, but construction decisions are rarely independent of physical context. The same issue can have a completely different impact depending on where it occurs, what surrounds it, which systems serve the area, and what work is scheduled next.

Impact depends on where

An AI system may recognize the words 'failed inspection,' but the project needs much more. Did the failure occur in a typical office, an electrical room, a rated shaft, an occupied healthcare area, or a ceiling plenum scheduled to close tomorrow? Does the condition appear once, repeat across similar rooms, or continue vertically through several levels? Which drawings, model elements, schedule activities, forms, and responsible parties are connected to that location?

If location is stored as inconsistent free text, AI must guess. Room names differ. Levels are numbered differently. Areas are duplicated. Systems use incompatible hierarchies. A model element may know its Revit room while an inspection knows only a field-entered description. The AI can produce fluent language, but it lacks a dependable spatial foundation for reasoning.

> *The AI can produce fluent language, but it lacks a dependable spatial foundation for reasoning.*

A governed where

blade ai's LBS provides that foundation by establishing a governed 'where' dimension. Locations have stable identities, controlled types, parent-child relationships, full codes, and complete paths. The client can use project-specific labels while the underlying structure remains consistent for queries, integrations, and analytics.

This allows AI to reason over location as structured knowledge rather than a collection of strings. A question about Level 4 can include all of its zones, rooms, corridors, cores, inspections, issues, forms, and connected evidence. A query about a turnover area can identify incomplete work across multiple categories. A recurring defect can be compared across similar units instead of being hidden behind small naming differences.

More than a tree

The hierarchy is only the beginning. LBS represents containment: where an object or record belongs. Other spatial relationships belong in a graph layer. A wall can bound two rooms. A corridor can be adjacent to a suite. A shaft can intersect multiple levels. A system can serve several spaces. Keeping these relationships outside the core hierarchy preserves a clean source of truth while allowing richer spatial intelligence above it.

blade ai's Vertical Zones address another reality that simple location trees miss: some construction conditions continue across floors. A riser, shaft, stair, or other vertical assembly cannot be understood correctly when it is duplicated as unrelated locations on each level. A relational cross-level grouping gives AI and users a way to reason about continuity without compromising the Primary LBS.

The same principle applies within a room or area. Vertical segmentation can distinguish full-height space, occupied space, ceiling plenum, and floor plenum. That context can materially change which trades, inspections, details, and risks apply. 'Room 201' is not enough when the question concerns above-ceiling MEP work versus finishes in the occupied zone.

Trust comes from the path

For clients, spatial intelligence can improve recommendations and prioritization. blade ai can help identify areas at risk of missing close-in, locations with incomplete evidence, repeated failures across similar units, downstream activities affected by an issue, or records that appear inconsistent with the governing location. These capabilities become more trustworthy because the system can show the spatial relationships supporting the conclusion.

A governed LBS also helps separate confidence from guesswork. When a record has a verified location assignment, blade ai can explain the path and related evidence. When the assignment is missing or ambiguous, the system can surface that limitation instead of pretending certainty. Human oversight remains central, especially where location affects safety, inspection, payment, schedule, or acceptance.

AI does not understand a construction project merely because it can read the documents. It must understand how the project is organized in space and how records relate to that spatial structure. LBS supplies the stable containment model. The graph and evidence layers add relationships and proof. Together, they allow blade ai to reason more like an experienced project team-starting with the essential question: where?

Read online: inertiasystems.com/resources/blog/why-construction-ai-must-understand-where

© 2026 Inertia Systems, Inc. All rights reserved.