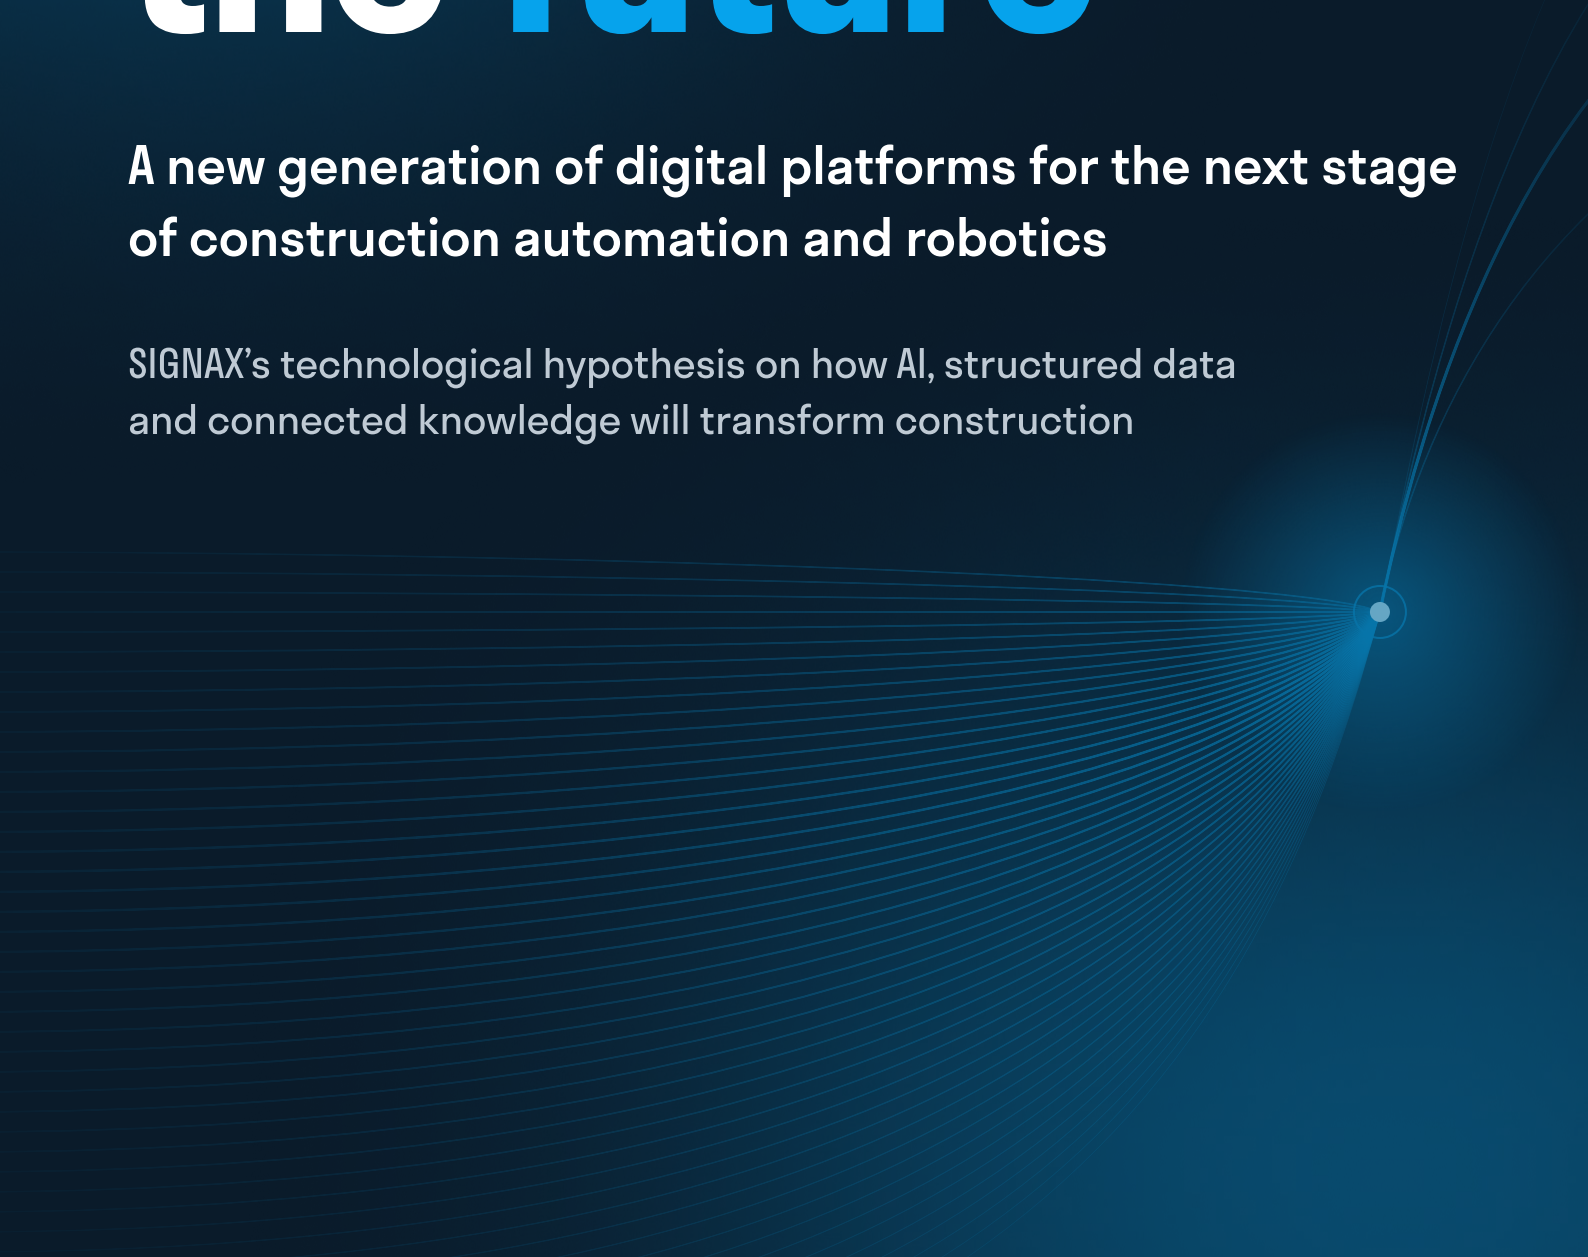
Turn and face the future

A new generation of digital platforms for the next stage of construction automation and robotics

SIGNAX's technological hypothesis on how AI, structured data and connected knowledge will transform construction



A SIGNAX White Paper

Alex Popov, Roman Sigitov 2026

Contents

	Foreword	2
	Executive Summary	4
01	Why Construction Has Hit a Ceiling	6
02	The Convergence of BIM, AI and Robots	9
03	The Evolution of CDEs and Their Architectural Limits	12
04	Three Architectural Shifts	16
05	Construction AI Harness: An Environment for Human–AI Collaboration	26
06	The Horizon: People, Agents and Robots	35
07	Why SIGNAX?	38
	Conclusion	43
A	Appendix A. Conducting an Experiment in Reviewing Design Documentation Using AI Models	45
B	Appendix B. Glossary	48
	About the Company	50

Foreword

Building the future: from digital construction to intelligent operations



The construction industry continues to face challenges in productivity, manpower, fragmented information, and increasing project complexity. While technology has helped us make progress, we still have much to improve in how we manage information, coordinate work, and make decisions.

Through my experience leading digital transformation at Boustead Projects and Boustead Singapore, I have learnt that successful transformation is not about adopting more technology. It is about solving real problems, improving the way we work and delivering measurable outcomes.

Over the years, Building Information Modelling (BIM), Integrated Digital Delivery (IDD) and Common Data Environments (CDEs) have provided an important foundation. However, information remains fragmented across systems, and many processes still rely heavily on manual effort. Artificial Intelligence (A.I) presents an opportunity to address some of these challenges, provided we apply it with a clear purpose.

This white paper by SIGNAX introduces the concept of a Construction AI Harness and its progression towards a Construction Operating System. It offers an interesting perspective on how project information, workflows and technology can be better connected to support construction delivery and operations.

What matters to me is how these capabilities translate into practical value. Can we reduce repetitive tasks, identify risks earlier, improve coordination and enable our teams to make better-informed decisions?

More importantly, Human Intelligence must remain at the centre of this transformation. Technology can process information and support decision-making, but it cannot replace engineering experience, professional judgement and accountability. The real value comes from combining human expertise with the right technology.

At Boustead, we have always focused on practical implementation, responsible adoption, and outcomes that make a difference to our people, projects, and business. Technology must serve a purpose, be scalable, and deliver value beyond initial implementation.

I commend SIGNAX for contributing to this important discussion and encouraging the industry to rethink how construction information and operations can be better connected.

Ultimately, the future of construction is not about how much technology we adopt, but how effectively we bring together Human Intelligence, processes, and technology to deliver better outcomes for the Built environment industry.

Muhammad Khalil Bin Shaiful

Senior Vice President, Group Technology
Head, Group Technology Office
Boustead Projects, Boustead Singapore

Executive Summary

Construction needs a new technological paradigm, rather than another specialized application.

1. **The industry has hit a ceiling.** Projects are becoming more complex, more parties are involved, a smaller workforce is available, and management still relies on email, messaging apps and spreadsheets. The root cause lies in data architecture: information is fragmented, changes go unchecked, and management systems cannot see the construction site.
2. **BIM, AI and robotics are converging.** They all need the same thing: a shared, machine-readable project context. The natural candidate for this role is the Common Data Environment (CDE), which already holds both project data and the industry's established business processes.
3. **Today's CDEs are reaching an architectural limit.** They manage documents and processes, but not the state of the physical asset. Adding an AI agent on top of documents is not enough.
4. **Three architectural shifts are needed**: from documents to project and company memory, from files to AI-native formats, and from processes to operations. These shifts turn the CDE into a Construction AI Harness: the environment in which AI works. The platform manages context, provides the model with specialized tools and performs deterministic operations, while the LLM reasons, orchestrates processes and interacts with people.
5. **Harness quality** is becoming as important as model quality. A good environment does not make the LLM read the entire project, write supporting code or perform calculations through token generation. It finds the required data, converts formats,

runs precise queries and calculations, and gives the model only the context it needs. As this environment begins to coordinate the state of the physical asset, people, agents, equipment and construction operations, the Construction AI Harness evolves **into a Construction Operating System**.

6. **The outcome** is a Construction Operating System (COS): a coordination layer connecting data, people, AI agents and, eventually, robots, while preserving human control over every critical action.
7. **The SIGNAX platform** brings together modular architecture, Harness layers, embedded AI agents and data sovereignty to unite information, knowledge, AI tools and machine-readable processes in a single environment for human–AI collaboration.

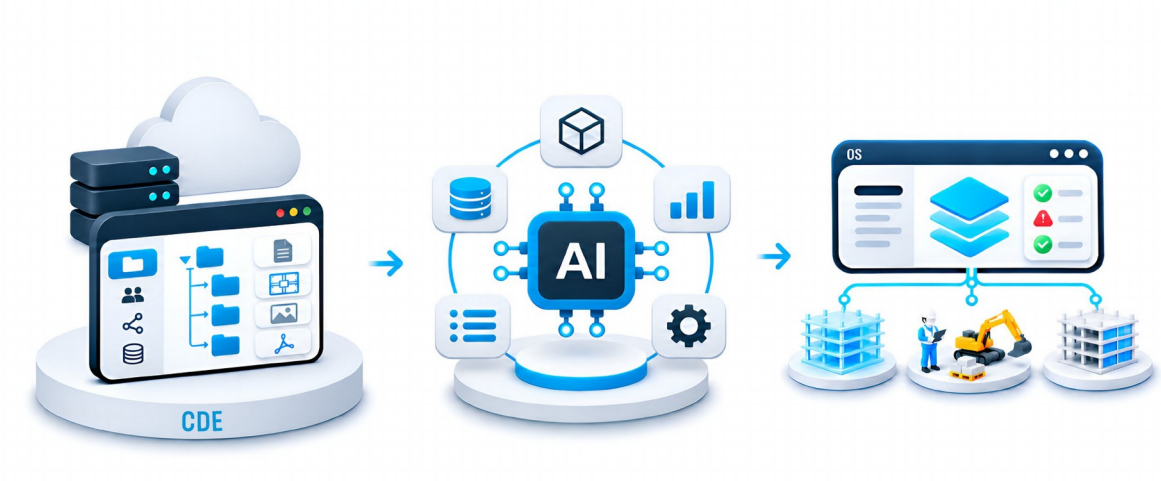


FIG. 1. **From CDE to a Construction Operating System**

The CDE organizes documents, entities and processes. The AI harness adds project and company memory, tools and LLMs. The Operating System coordinates the physical state of the asset, people, agents and equipment.

1. Why Construction Has Hit a Ceiling

Construction is the largest sector of the global economy and also one of the least efficient. According to the McKinsey Global Institute, construction labor productivity grew by an average of 1% annually over two decades, compared with 2.8% across the global economy, while large projects typically finish 20% behind schedule and up to 80% over budget.[1] Engineers and managers spend a substantial share of their time searching for information, obtaining approvals and coordinating with other disciplines, rather than on production tasks.

Logistics, retail and finance automated their core processes long ago. Construction still relies on analog tools and paper processes that are manually transferred into digital systems and back again. The problem is not a lack of effort: three global trends have converged in construction at the same time.

1.1 Three Forces

Projects are becoming more complex. Safety, environmental and energy-efficiency requirements are multiplying building MEP systems. Modern architectural forms demand complex structural and envelope systems, normal-weight concrete and high-strength steel. Schedules remain tight, so cast-in-place concrete, facade and MEP works proceed simultaneously on different levels, sharply increasing the demands on coordination and workforce skills.

More parties are involved. A large project brings together thousands of workers and hundreds of engineers from dozens of contractors, suppliers and consultants, often from

1 McKinsey Global Institute, "Reinventing Construction: A Route to Higher Productivity," 2017.

different countries. Each organization issues hundreds of requests, emails and documents recording continuous changes over many years, from site investigations through commissioning. Since the Tower of Babel, the greatest threat to a large construction project has been a breakdown in communication.

Smaller workforce available. Demographic trends are reducing labor supply in the United States, Europe, the Gulf states and Southeast Asia. Engineers are also in short supply: in its current form, construction is losing to other industries in the competition for talented graduates.

1.2 The Low-Margin Trap

Increasing project complexity and workforce shortages increase costs, while tendering procedures and the transfer of risk to contractors squeeze already thin margins. Little money remains for digitalization. As a result, most companies rely on email, messaging apps, spreadsheets and file storage as their main management tools: disconnected systems whose architecture cannot match the complexity of the work. Market leaders use advanced systems, but this only highlights the gap with the rest of the industry.

1.3 The Root Cause Is Data Architecture

As a result, most companies operate within a digital environment in which:

- **data is fragmented** into separate, disconnected documents; BIM has yet to become a unified project information system across mass construction;
- **information is scattered** across formats, stakeholders, folders, attachments, correspondence and paper records on site;
- **changes do not propagate**: no system automatically checks them or carries them through the chain of dependencies, so information is lost or distorted at every step;
- **management systems are blind**: ERP records transactions and resources but cannot see the physical state of the asset.



FIG. 2. **Why today's systems cannot see the construction site**

Data is fragmented, information is scattered, changes do not propagate, and management systems cannot see the physical state of the asset.

"No people, no profit, no accessible management tools": this is what the construction industry is saying today.

Another application will not solve this problem. A different architecture for data and management is needed.

2. The Convergence of BIM, AI and Robots

2.1 Promising Technologies

The industry is undergoing two technological transitions at once, with a third approaching:

1. **BIM**: digital information models of built assets. For now, these are used mainly in design and rarely reach the construction site.
2. **AI**: large language models capable of organizing vast amounts of data and solving problems across disciplines.
3. **Robotics**: unmanned aerial vehicles (UAVs), autonomous machinery and, eventually, humanoid robots.

These technologies developed independently, but they all need the same thing: **a shared, machine-readable project context** covering what is intended, what has been designed, what has been built and what is happening right now. We believe the adoption of robotics in construction is inevitable, and the transition from building by hand to building with autonomous machines will require a new digital environment.

2.2 Our Outlook: Construction behind IT by 3–5 Years

Our method is simple: construction digitalization follows the path of software development with a lag of approximately 3–5 years. The way software and game developers use AI today indicates what will emerge in design and construction once the industry has tools of comparable maturity. We draw four lessons from IT that will underpin this technological transition:

- **Context is a limited resource.** Developers have learned to manage what is loaded into the model: when to reset context, when to compress it, and how to divide data so that only the necessary parts are retrieved.
- **Algorithms wherever possible; the model wherever necessary.** Modern AI development tools operate as a **harness**: a software environment around the LLM. It gives the model tools to read and modify data, search, convert formats, calculate and interact with other systems. The environment performs deterministic operations, while the model understands the task, selects tools, orchestrates their calls and reasons. This approach improves output quality, accelerates work and reduces token consumption at the same time.
- **Specification, execution, acceptance.** In spec-driven development (SDD), a person defines the task, an agent executes it, and the human accepts the result.
- **Open ecosystems win over "all-in-one" solutions.** In our case, this means a family of standardized AI-native representations of construction data. See Section 4.2 for details.

2.3 Why the CDE Becomes the Hotspot

Information from every participant—developers, designers, contractors and suppliers—passes through the Common Data Environment: design, construction and as-built documentation, BIM models, issues and approvals, site instructions and checklists, logs, schedules, reports, site photos and videos, correspondence and management metrics. In effect, the CDE forms a digital reflection of the life cycles of many assets: a context that cannot be recreated artificially.

But the CDE holds more than data. **Its configuration captures the industry's working business processes.** Entity types, their relationships, approval workflows, participant roles and permissions form a software model of how a company actually designs and builds. There is no better description of construction business logic: this is not a procedure someone wrote on paper, but logic tested every day through real work.

The CDE is therefore the natural core of the future Construction AI Harness. It already contains project data, roles, permissions and most of construction's real business processes.



FIG. 3. **A single machine-readable context for BIM, AI and robots**

Information from all participants flows through the CDE: what was intended, what was designed, what has been built and what is happening right now. This context is needed by people, AI agents and, in future, equipment and robots.

However, connecting a powerful LLM to an existing CDE is not enough. The environment around it must be designed: which data enters context, which operations are performed by algorithms, which tools the model can use, and how its outputs are verified. The quality of this environment will largely determine the quality of the applied AI system.

Once that environment begins to coordinate not only digital processes but also the physical state of the asset and construction operations, it can evolve into a **Construction Operating System**.

3. The Evolution of CDEs and Their Architectural Limits

CDE Generation	Primary Object	Who Acts	The Question the System Answers
1. Document-centric — yesterday	Files and documents	A person	“Where is the current version?”
2. Workflow-centric — today	Processes and entities	A person, guided by an AI assistant	“What stage has the approval reached?”
3. Harness-centric — tomorrow	Knowledge and processes	People and AI agents under human supervision	“What is happening, and what is likely to happen?”
4. Operation-centric — the day after tomorrow	Operations and the state of the asset	AI agents	“Who should do what, where and when next?”



FIG. 4. Four generations of CDEs

Documents → processes and entities → knowledge and processes → operations and asset state. The question the system answers changes from “Where is the current version?” to “Who should do what, where and when next?”

3.1 Generation 1: Document Storage

First-generation CDEs, such as Aconex, which appeared more than 20 years ago, brought order to project document management: storing and transmitting drawings and documents, legally significant correspondence, access rights and file versions. Documentation moved through these systems from the client to the designer, from the designer to the contractor, and back to the client and the asset operator.

The early versions of these systems were not designed to work with BIM models and had no mobile applications.

3.2 Generation 2: Creating and Managing Processes

Over the past decade, second-generation CDE leaders Autodesk Construction Cloud (ACC) and Procore have moved from storage to process management: routed document approvals, requests for information (RFIs) and issues, work planning, acceptance and certification of completed work, reporting, smartphone inspections, and budgets and cost changes linked to contracts. ACC gained native integration with Revit, Navisworks and AutoCAD: model collaboration, clash detection and issues linked to elements.

The key change is that **the document has given way to the entity**. An RFI, site instruction, NCR or acceptance certificate is described by a set of attributes: originator, recipient, BIM element, site location, priority, deadline, status, related documents and decisions made. The system begins to work with objects in the production process, rather than files.

In 2026, the second generation became agentic. In June, Procore introduced more than 150 agentic actions for creating and modifying RFIs, Submittals, Daily Logs and other objects. In September, Autodesk announced AI agents operating on project data and positioned ACC as an end-to-end AI-native AECO platform.

3.3 Generation 3: Knowledge Management

Second-generation agentic features accelerate work with documents and processes. But an architecture designed for people limits what AI can do:

1. **The agent operates on top of documents rather than within the project model.** To understand the context, it rereads the files every time. This is expensive, slow and does not scale to the entire project.
2. **Formats are not designed for AI.** Proprietary CAD formats are accessible only through APIs, while open IFC relies on cross-references between lines and was created as an exchange format, rather than a format the model can generate and edit (Section 4.2).
3. **The process does not know the state of the asset.** A workflow records that a document has been approved, but not that a structure has been constructed, accepted and paid for.
4. **There is no execution and verification loop.** The system does not assign tasks to people on site or confirm physical results.
5. **There is no memory or context management.** Project and company knowledge does not accumulate in a form convenient for the model, and every request starts from scratch.
6. **There is no harness infrastructure.** Agent permissions, context management, available tools, human confirmation, action auditing, token control and emergency shutdown are not built into the core of existing CDEs.

These limitations cannot be removed by adding another AI agent on top of an existing CDE. An architectural transition is needed: the CDE itself must become an environment optimized for collaboration between people, algorithms and LLMs—a **Construction AI Harness**.

Documents, BIM models, schedules and cost estimates become sources of knowledge for AI agents, rather than the final outputs of a digital process. Using that knowledge, agents will be able to generate specific operational tasks—deliver materials, carry out installation, perform an inspection, resolve an issue, accept completed work—and support the full construction cycle: brief → design solution → production task → physical operation → result verification. Eventually, robotic machinery will carry out these tasks alongside people and companies.

3.4 Generation 4: Operations Management

Finally, the next stage of CDE development is the transition from managing project knowledge and information processes to a machine-readable operational model of the project. The CDE will then evolve into a **Construction Operating System (COS)**: a platform connecting project data with real actions, generating and initiating decisions, assigning tasks to designers, contractors and suppliers, coordinating work, and verifying and accepting the result.

4. Three Architectural Shifts

4.1 From Documents to Knowledge: Project and Company Memory

The previous stage of digitalization turned paper into digital information. The next turns accumulated data into knowledge accessible to AI. A platform's value is determined less by its feature set than by how quickly and accurately it gives the model the context it needs.

Memory matters more than entity descriptions. Through training, the model already knows what a site instruction, contractor or acceptance certificate is and how they relate. It learns the rest from the CDE structure: there are relatively few entity types, with relationships defined by fields and links. The main task is therefore to organize continuously updated project, company and external context memory from which the model can quickly retrieve what it needs.

The processing pipeline. Every document entering the CDE is automatically:

1. recognized: scanned and PDF images are converted to machine-readable digital text;
2. parsed into structured data and fragments of a manageable size;
3. indexed and classified;
4. vectorized and stored in a vector database;
5. assigned to the appropriate memory level.



FIG. 5. **How a document becomes project memory**

A document is recognized, parsed, indexed and vectorized, then assigned to a memory level. The harness works with this structure and retrieves only the information the model needs.

Three memory levels.

Level	What Is Stored	Purpose
Project memory	Technical briefs, explanatory notes, documentation sets and their contents, project contracts, participant roles and obligations, and data about the physical asset	Context for tasks within a specific project
Company memory	Standards and templates, standard solutions, contractual obligations and deadlines, counterparty profiles and working history, and recurring risks	End-to-end analytics; the experience of past projects influences current ones
External context memory	Regulations, public reference sources and market data	External constraints and practice

Which data is stored in the company's memory and which remains at the project level are determined by the configurable rules of the platform. For example, a contractor moves from project to project, so its profile and work results remain in the history of the company's memory, while its role and obligations under a specific contract belong in the project memory.

Sources carry different weights: some documents determine decisions, while others provide supporting information. In software development, project memory often fits into a

single Markdown file; a construction project needs a vector database with configurable source importance.

Context management. The model's context window is limited, so the system:

- divides large documents into entities and fragments so that only the necessary parts are retrieved;
- searches the vector database for relevant fragments when project-wide or company-wide context is needed, rather than loading everything;
- answers precise questions—"how many NCRs are open on the third floor?" or "which acceptance certificate covers this structure?"—by querying structured data rather than the model;
- compresses and resets context in long sessions.

This is one of the fundamental functions of a Construction AI Harness. Its task is to assemble **the minimum sufficient context** for a specific operation, rather than pass as much information as possible to the model. The less irrelevant data occupies the context window, the more capacity the model retains for analysis, reasoning and producing a result.

What this enables. Company memory makes it possible to assess a project against a particular company's standards, past decisions, preferred materials, recurring risks and procedures, rather than average industry practice. This is especially important in highly specialized segments, where general market patterns differ from internal practice. Eventually, anonymized knowledge from different companies could be combined into industry models, but only with the consent of the data owners (see Section 7.2).

4.2 From Files to AI-Native Formats

Proprietary formats—Revit, Civil 3D and Tekla—are closed: AI can work with them only through programming interfaces, rather than directly with the data.

Why IFC Is Unsuitable for AI-Driven Design

IFC is the principal open format for exchanging BIM data. But the future of AI-driven design lies in a model that not only analyzes a design but also generates and edits it, just as AI writes software code today. IFC has fundamental limitations for this task.

1. **The meaning of the object is scattered across the file.** In the IFC text representation, each record refers to others by line number:

```
#123=IFCWALL(...); #124=IFCMATERIALLAYER(...);
#125=IFCRELASSOCIATESMATERIAL(...,#123,#124);
```

To understand information about a single wall, records from different parts of the file must be assembled. This is inconvenient for people, and even more so for models.

2. **A small edit triggers a cascade of changes.** During export, optimization, model merging or resaving, many IFC generators rebuild the file and reassign identifiers. A change to a small fragment affects thousands of lines, even when the project's geometry and logic have barely changed.
3. **IFC is poorly suited to Git.** In software development, Git stores history, shows differences between versions, merges contributors' edits and provides “before → after” datasets for model training. This requires small changes to affect only a few lines. IFC behaves in the opposite way: diffs are unreadable, merges cause conflicts, and the change history is almost impossible to analyze.

The model inherits the format's limitations. A neural network trained on IFC spends resources not on building design, but on thousands of interconnected STEP records, service relationships, and format syntax. It is like writing a program directly in machine code instead of a high-level language.

2D documentation is left out. IFC only partially describes plans, sections, elevations, nodes, electrical, single-line and technological schemes, dimensions and annotations. As a result, the model is stored in IFC, drawings in DWG, and diagrams in other formats, with their relationships maintained separately. Training AI therefore requires collecting data from many disconnected sources.

IFC remains a good transit format between applications, but not the primary carrier of project knowledge. AI-driven design needs a format in which:

- each element is described independently;
- stable object GUIDs replace positional line identifiers;
- changes are local, and Git works as it does with code;
- the 3D model, 2D drawings, grids, dimensions, annotations and schedules are described within one system;
- the structure is understandable to both people and neural networks.

Such a format should play the same role in construction that source code plays in programming.

The Markdown lesson. Markdown was created long before LLMs but proved an ideal format for human–AI collaboration: it can be read without specialized software, is easy to generate, contains almost no syntactic noise and divides naturally into logical blocks. An entire ecosystem has grown around it, including Notion, Obsidian and Logseq. The lesson is that simple, structured, human-readable representations work equally well for people, software and models.

A family of domain-specific notations. The next step is specialized formats for construction entities. A technical brief contains blocks such as @PROJECT, @CUSTOMER, @REQUIREMENTS, @CONSTRAINTS and @DELIVERABLES; a contract uses @PARTIES, @SUBJECT, @PAYMENTS and @RESPONSIBILITIES; a schedule is described by tasks, dependencies and milestones; a budget by entities, calculation rules and computed fields. The model understands these documents through predefined constructs rather than natural-language parsing.

The structured description becomes the primary source, while each data type has its own Viewer, Editor and Validator for people: a 3D viewer for models, a CAD visualizer for drawings, an interactive chart for schedules, a financial form for budgets and inspection forms for quality control.

Each element is a separate entity. In an AI-native BIM format, each element is described by a self-contained record with a stable identifier, and relationships use simple references rather than deeply nested structures:

```
WALL;ID=W-0412;L2;X1000;Y2000;X5000;Y2000;T200;TYPE=W01 DOOR;ID=D-0107;L2;HOST=W-
0412;X3000;Y2000;W900;H2100;TYPE=D01 ROOM;ID=R-0203;L2;NAME=OFFICE;AREA=24.5
```

This approach enables streaming generation, independent object editing, work within limited context, simple change validation and convenient storage in version control systems.

Decomposition. The asset is divided simultaneously by discipline, system, floor, zone, work package and stage. One enormous file gives way to many connected parts, and the model loads only the part needed for the task:

```
ARCH/ L01/ L02/ L03/ STRUCT/ L01/ L02/ L03/

MEP/ HVAC/ ELECTRICAL/ WATER/
```

SIGNAX's SG Format

The SG format is our response to these requirements: a family of standardized AI-native representations of construction data.

- **JSON as the foundation.** A standardized schema is convenient for transmission through REST APIs and web interfaces. For input to the model, we are also considering a compact Markdown representation that saves tokens. The final balance between the two is currently being determined.

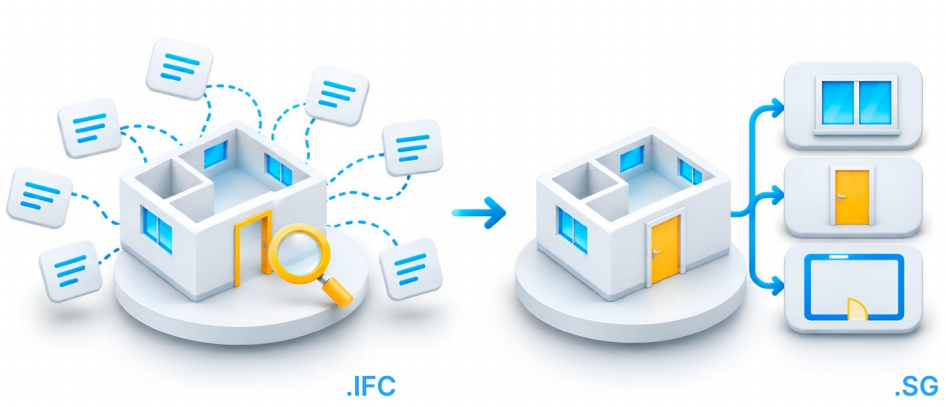


FIG. 6. **How IFC and SG represent the same element**

In IFC the meaning of an object is spread across linked records; in the SG format each element is a self-contained record with its own attributes and relations, with a stable ID, so edits stay local and a model can load only what it needs.

Instant indexed database. Structured JSON is quickly reassembled into a local database: the properties of a medium-sized BIM model, around 300 MB, are converted in about 300 ms—directly in the browser. Compared with chat response times, this delay is imperceptible to the user.

Algorithms instead of tokens. This database supports fast queries, calculations, validation and aggregation: quantity takeoffs from a BIM model, checks that required parameters are populated, and parameter selection for search sets. Algorithms perform this work, rather than the model (Section 4.3). Within the Construction AI Harness architecture, these operations become specialized tools that the LLM can call when needed. The model does not have to load the source dataset into context or write a program to process it every time: it formulates a tool request and receives a compact result.

Formats of the family:

Format	What It Describes	How AI Uses It
SGBIM + SGPROPS	BIM model geometry and properties in two files; together, the unified model storage format in S BIM	Queries properties without loading geometry; already used in SIGNAX
SGMPP	Construction schedule	Creates and adjusts the schedule through voice or text requests in chat
Text documents	Technical briefs, explanatory notes and contracts	Generates documents, regenerates a selected fragment following user instructions, and compares documents

How SG differs from the alternatives:

Criterion	IFC	Arbitrary JSON	SG format
Ability to be generated and edited by the model	Difficult: object meaning is distributed across linked records	Possible, but without a unified schema	Domain-specific schemas; one record per element subject diagrams, an element is a separate entry
Locality of changes and Git compatibility	A small edit changes thousands of lines	Depends on the structure	Changing an element affects only its record
Partial loading into context	Difficult	Depends on the structure	Decomposition is built into the format
API transmission	Large files	Yes	Yes
Precise queries and calculations	Requires a specialized engine	Requires parsing	Local indexed database

Two Stages: First, MCP on Top of CAD, Then AI-Native CAD

IT experience suggests how the market for design tools will change. Replit was conceived as an AI-first development environment, and many expected it to replace VS Code and other IDEs. But then plugins, MCP servers and agents configuring CI/CD in GitLab appeared, giving traditional IDEs AI capabilities while remaining more powerful, cheaper and more convenient for manual code editing and large projects. Many developers returned to them. We expect the same in building design:

- **Over the next 2–3 years**, development will center on Revit, AutoCAD and other mature CAD systems through plugins and MCP servers. For now, AI assists the designer with routine and targeted operations, rather than generating an entire design. The designer needs a full-featured tool to edit the result manually and assign the model a specific task directly in the open file.
- **Later**, AI-native CAD systems may replace today's tools by operating within a unified system: a complete version history, branches, commits and merges, as in Git; clear visualization of changes; and end-to-end design following SDD, where a task is defined, an agent executes it, and the designer accepts the result.

The SG format is needed at both stages: today as an AI-accessible representation of data from existing CAD systems, and tomorrow as the foundation for models created and edited by agents.

4.3 From Processes to Operations

Machine-readable processes. When a process is described as a sequence of operations, transition conditions, roles, permitted scenarios, deadlines, business rules and automated checks, both people and AI agents can understand it. The agent analyzes progress, identifies deviations, assigns tasks, monitors deadlines and coordinates participants.

Events replace manual monitoring. Agents are activated by events: a document is approved, an RFI is answered, a BIM model changes, a check fails, a deadline passes, materials arrive or a payment is confirmed.

The CDE becomes a Construction AI Harness. As the application environment around the model, the Harness stores and prepares data, manages context, provides specialized tools, performs precise operations and controls AI actions. Token savings become a key performance metric for this platform: the more work algorithms can reliably perform, and the more precisely the system selects context, the more efficiently the LLM works:

Actor	Responsibilities
Platform algorithms	Structured data queries, calculations, rule-based checks, process routing, aggregation and deadline monitoring
LLM	Understanding user intent and unstructured input; reasoning; selecting and calling tools; orchestrating processes; analyzing ambiguous situations; interacting with humans
People	Goals and priorities, decisions, confirmation of critical actions and acceptance of results

The operational cycle. In an operation-centric system, every action follows the same closed loop:

data → validation → decision → task → execution → verification → knowledge update

Documents become data. Future processes should not be designed simply by transferring today's paper forms into digital form. When a process is fully digitalized, many documents just disappear altogether. An American truck driver already operates within a digital process: Rate Confirmation defines the trip's terms, the BOL accompanies the cargo, and Proof of Delivery confirms the result, with data passing sequentially between the broker, carrier, driver, recipient and payment systems. People perform physical operations while the information process increasingly exists and moves independently of them. Likewise, a car-sharing customer does not sign a paper rental agreement: there is a digital record exchanged between systems. Digital government services such as Singapore's Singpass and the UAE's UAE Pass show how bureaucracy shrinks simply through digitalization. Singapore's CORENET X already accepts design documentation for approval as a BIM model.

The same will soon happen in construction. Work quality will be confirmed through a data query rather than a separate certificate. The remaining documents will not replicate today's one-to-one forms: fewer fields, links to common data accessible to both parties, and no duplication. During the transition, a paper version will be a separate process: export the current state and generate the form required by the authority. Ultimately, exchanges with authorities will also take place through APIs.

5. Construction AI Harness: An Environment for Human–AI Collaboration

A large language model alone is not yet an applied AI system.

An LLM can understand a task, work with unstructured information, reason, select the next step and interact with humans. But to do real work, it needs a software environment around the model: access to data, memory, tools, permissions, processes, result verification mechanisms and ways to interact with external systems.

This environment is called a Harness.

The Harness determines both what the model can do and how efficiently it can do it.

A good harness does not make the LLM perform work that a conventional algorithm can do faster, more cheaply and more reliably. The model should not have to write a file-conversion script every time, parse a proprietary format on its own, calculate thousands of values or load hundreds of documentation pages to find a single parameter.

Instead, such an environment provides the model with ready-made tools. The LLM understands what needs to be done, selects the appropriate tool and interprets the result.

In construction, this environment is the Construction AI Harness.

5.1 Harness Quality Becomes as Important as Model Quality

Two systems using the same LLM can produce entirely different results.

The first gives the model hundreds of documentation pages and asks it to find errors.

The second first defines the task, finds the necessary documents, extracts only relevant data, runs formal checks programmatically, provides compact context, lets the model request additional data through tools, and then programmatically rechecks part of the result.

In the second case, the same model works faster, more cheaply and often more accurately.

The quality of an applied AI system is therefore increasingly determined by the combination of two components: **Model Intelligence × Harness Quality.**

The model determines reasoning quality. The Harness determines the quality of data, context, tools and execution.

5.2 Context Is a Limited Resource

A construction project can contain thousands of documents, BIM models, contracts, schedules, correspondence, photographs, issues, survey results, regulations and years of change history.

Loading all of this into an LLM is neither possible nor necessary.

Even as context windows grow, context remains a limited resource. The more irrelevant information it contains, the more expensive the request, the greater the latency, and the harder it is for the model to identify the data that actually matters.

Context management is therefore one of the key functions of a Construction AI Harness.

The system must determine:

- what the model needs to know for a specific task;
- what should be retrieved through search;
- what can be queried from a structured database;
- what should be calculated programmatically;
- which data needs to be converted;
- and what should not enter context at all.

The Harness's task is to assemble the minimum sufficient context for the model.

This both reduces token consumption and improves model performance.

5.3 The CDE Becomes a Construction AI Harness

Today, a CDE primarily answers the question: **“Where is the information I need?”**

A Construction AI Harness must answer a different question: **“What does AI need to know, and which tool must it call, to perform the next operation?”**

The Common Data Environment already contains most of the necessary foundation: project participants, roles and permissions, documents and their versions, approval routes, BIM models, forms, issues, tasks, schedules, as-built documentation and actual site data.

But information must be organized differently for human–AI collaboration.

AI should not have to reconstruct the project structure from folder and file names every time. The Harness should already know where to find the technical brief, engineering surveys, initial data and permits, design solutions, construction documentation, schedule, contracts, issues and evidence of completed work.

This knowledge can be physically stored in different ways:

- in a structured database;
- in Markdown;
- in JSON;
- in SG formats;
- in conventional documents;
- in a vector database;
- in a relationship graph.

What matters is a stable, machine-readable structure understood consistently by the platform, people and AI, rather than the particular storage system.

The LLM should not have to work out which “technical_brief_final_final_3.pdf” file is current. The Harness should provide an unambiguous entity: the current project technical brief.

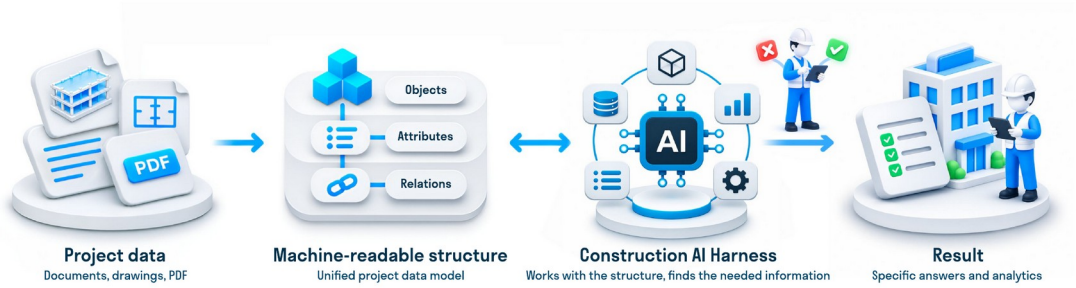


FIG. 7. **Construction AI Harness architecture**

The harness sits between project data and the LLM. It manages memory and context, provides tools and checks results, while people keep control over critical actions.

5.4 Documents Become an Interface for Human–AI Collaboration

A Construction AI Harness also changes how information is created.

A technical brief, for example, need not start with a blank Word file. The user starts the "Create a Technical Brief" process. AI collects the necessary information step by step: the asset's purpose, initial data, constraints, client requirements, deliverables and other parameters. The Harness checks that all required data is present. The LLM drafts the document. A human edits and approves it. Once approved, the structured brief automatically becomes part of project memory and a source of context for subsequent processes.

Another AI agent no longer asks the user about project requirements again: it knows where to find the current version.

The same principle can be used to create design briefs, decision records, RFIs, schedules, work plans, reports, certificates, specifications, contractor assignments and inspection checklists.

The document gradually changes from a separate file into a structured object for human–AI collaboration.

5.5 Workflows Become Executable Scenarios for Humans and AI

In a conventional CDE, a workflow primarily defines a document's route between people.

In a Construction AI Harness, the process becomes a machine-readable scenario whose participants can be: human → AI → algorithm → AI → human.

For example, “Create a Technical Brief”: a human initiates the process → AI collects requirements → the Harness checks required data → AI generates the document → the human edits it → AI checks consistency → the person approves it → the document is recorded in project memory.

Another example, “Review Construction Documentation”: a new version appears → the Harness identifies changed parts → algorithms extract structured data → AI analyzes changes → formal checks run programmatically → potential problems are rechecked → an engineer receives meaningful findings → confirmed findings become project entities.

This kind of workflow turns an isolated conversation with a neural network into a repeatable production process.

5.6 Construction AI Harness Tools

A Construction AI Harness provides the LLM with specialized tools.

- **Data tools.** Retrieve precise data without loading source files into context.
- **Search tools.** Find relevant information in project, company and external context memory.
- **Conversion tools.** Convert PDF, DWG, IFC, Revit and other formats into representations suitable for AI.
- **Calculation tools.** Perform calculations, queries and aggregation.
- **Validation tools.** Check formal rules.
- **Authoring tools.** Create and modify documents, schedules, models and other entities.
- **Workflow tools.** Start processes, create tasks and change states.
- **Integration tools.** Connect the Harness with Revit, AutoCAD, ERP, scheduling systems, IoT and other applications.

Some tools may be built directly into the platform; others may connect through APIs, plugins or MCP servers.

The guiding principle: do not use an LLM for work that an algorithm can reliably perform.

If 20,000 values need to be added, a program does it. If all doors of a certain type need to be found, this is a structured data query. If differences between two model versions need to be identified, a diff algorithm does it. If required parameters must be checked for completeness, a validator does it.

The LLM is used where meaning must be understood, ambiguity interpreted, a sequence of actions selected, or communication with a human is needed.

5.7 Practical Benefits

Documentation review. The quality of an AI review depends on the Harness as well as the model. The documentation set is analyzed systematically; data across documents is linked; numbers are extracted and compared; precise calculations are performed programmatically; cross-references are checked systematically; regulatory references are verified against controlled sources; each finding is linked to a location in the document; and potential findings are rechecked.

SIGNAX experiment, September 2026. The Structural Design section for a 21,600 m² industrial building, comprising 40 drawing sheets. The best model found 13 of 14 reference errors; an engineer confirmed 21 of 23 findings. The review required 10 minutes of model work and 25 minutes of engineer validation, compared with 220 minutes of manual review. Engineer time savings: 88.6%. See Appendix A for details.

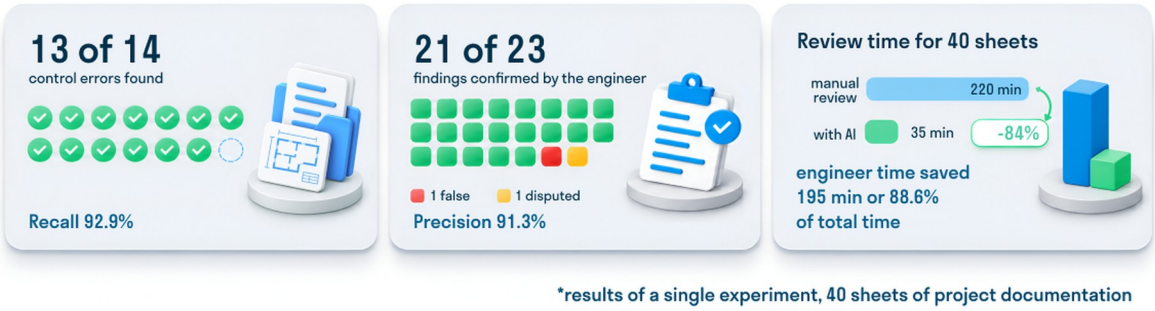


FIG. 8. **SIGNAX experiment results**

The best model found 13 of 14 reference errors; the engineer confirmed 21 of 23 findings. Review took 35 minutes instead of 220 minutes of manual work.

Early risk warnings. The Harness can compare the current project with accumulated experience and provide early warnings of patterns that previously led to schedule delays, growing backlogs of issues, resource shortages or declining work rates. This requires addressing data provenance and reliability, project standardization and the identification of truly comparable assets.

Generating solutions. To generate high-quality BIM models, schedules and other outputs, AI needs patterns from real projects, company standards and a regulatory knowledge base. The Construction AI Harness becomes the point where these sources are combined and provided to the model in a controlled form.

Agentic project management. An agent working within the Harness can explain why a stage is delayed or which work is affected by a design change, and can also act under human supervision: prepare an RFI, assign a task, start a process, escalate a delay and check that an issue has been resolved.

5.8 People Remain in the Loop

A Construction AI Harness does not mean handing a construction project over to a neural network.

Roles are divided: algorithms perform deterministic operations; the LLM understands, reasons and orchestrates; humans define goals, make decisions and confirm critical actions.

Therefore, the Harness must include:

- AI agent permissions;
- restrictions on available tools;
- confirmation of critical actions;
- an audit trail of all actions;
- versioning and rollback;
- token and cost control;
- detection of anomalous behavior;
- emergency shutdown;
- human verification of results;
- a choice between cloud and on-premises environments.

5.9 From Construction AI Harness to Construction Operating System

A Construction AI Harness primarily organizes collaboration between humans, AI and digital tools.

But as the system begins to know the state of the physical asset as well as the state of information, its boundaries expand.

If the platform knows:

- what has been designed;
- what has been approved;
- what has been purchased;
- what has been delivered;
- what has been installed;

- what has been inspected;
- what has been accepted;
- what has been paid for;

and can use this knowledge to coordinate the next actions of people, agents, equipment and robots, the Harness becomes a Construction Operating System.

The Construction AI Harness is the architecture of an AI-native construction environment.

The Construction Operating System is the stage of its development at which this environment begins to coordinate physical construction operations.

6. The Horizon: People, Agents and Robots

6.1 Robots on Site

Leading manufacturers are already developing autonomous robots capable of navigating space, moving over difficult surfaces and performing physical operations. This is still an early stage, but robots are highly likely to become full participants in construction within 10–20 years, adapted to its conditions: dirt and dust, temperature fluctuations, open voids, debris and unsecured structures. They will first take on supporting tasks such as delivering materials and tools, followed by discrete activities such as welding, installing partitions or laying blocks.

This raises a key question: **how will robots know exactly what to do?**

6.2 BIM as Spatial Context

> Today, a BIM model is primarily a 3D representation of the design and a tool for clash detection and quantity takeoffs; it is rarely used to assign tasks at the frontline production level. In the future, it will become part of robots' spatial positioning and coordination system, alongside SLAM, LiDAR and UWB. Using the BIM model, a robot will find a room, plan a route, locate the work area and compare the actual state of the asset with the design. BIM provides spatial and design context, while the Construction AI Harness connects it to operational context: installation sequences, quality requirements, constraints, tasks and work instructions. At a higher level of autonomy, when the same environment directly coordinates physical operations, it performs the functions of a Construction Operating System.

6.3 The Orchestration Hierarchy

The construction site of the future is a multilevel system:

- **robot workers** carry out local physical tasks;
- **specialized agents** monitor particular aspects: a quality agent analyzes video feeds and images from robot cameras, a schedule agent monitors timelines, and a cost agent monitors the budget;
- **orchestrator agents** at different levels distribute tasks, check execution and delegate to the level below, from local operations in a work zone to project-wide processes.

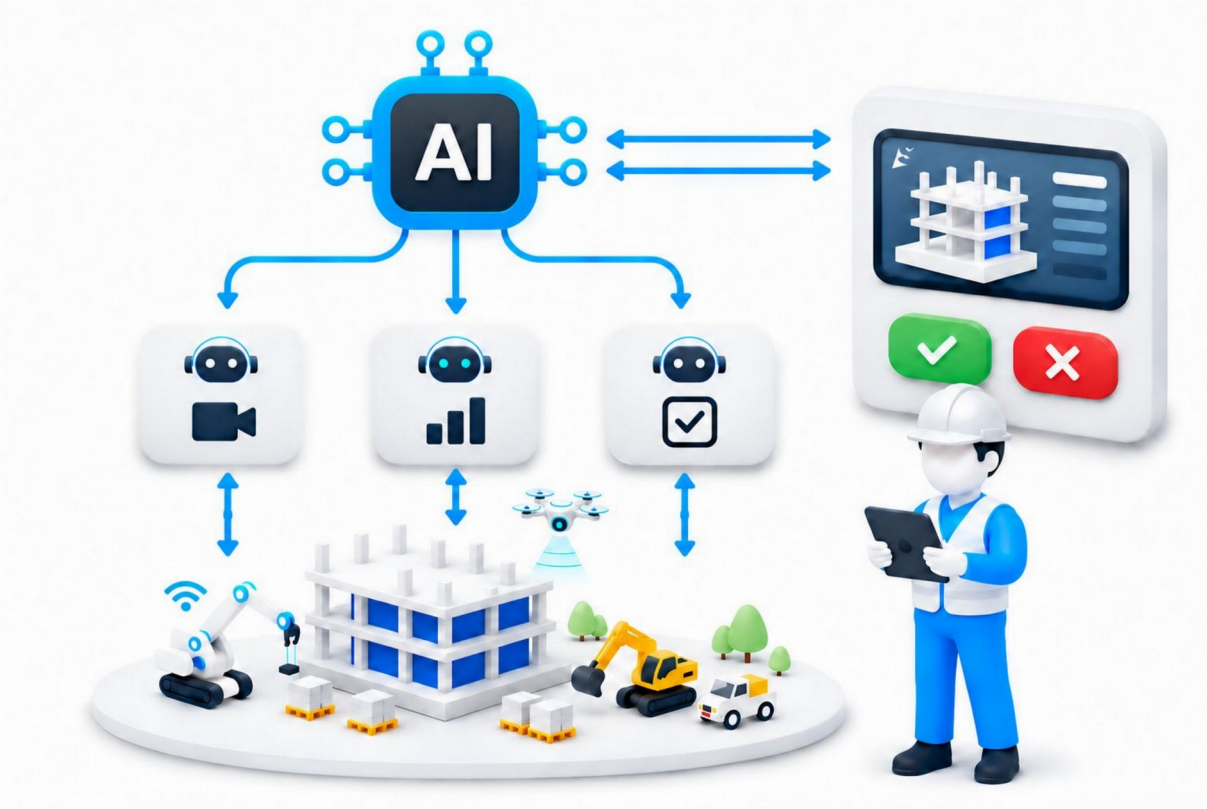


FIG. 9. **Who manages agents and robots**

Robot workers, specialized agents and orchestrator agents form a multilevel hierarchy, with people supervising at every level.

The business process structure may remain similar to today's. Agents will take the places where human managers currently relay information: first, an agent supervising individual activities—a “foreman”; then, a construction site coordination agent—a “site manager”;

then, a project management agent—a “project manager.” They will manage other agents and, ultimately, robots and autonomous equipment.

6.4 The Autonomy Ladder

Human involvement will decrease gradually:

Level	What the System Does	The Human Role
1. Assistant	Reviews documents, finds errors and offers guidance	Performs all work and makes all decisions
2. Supervised executor	Prepares solutions, assigns tasks to people and starts processes	Confirms every action
3. Dispatcher	Distributes tasks among people, robots and equipment	Handles exceptions and accepts results
4. Site coordinator	Agents manage agents and robots; robots recharge themselves, and some robots service others	Maintains an overview of the project and intervenes when necessary

At the early stages, humans participate in almost every process as controllers, observers and reviewers, and sometimes work alongside robots. As trust grows, they grant the system more autonomy and supervise at progressively higher levels. Eventually, one person can view the entire project on a single situational dashboard, understand what is happening through high-level analytics, and switch to an individual robot at any time. **Humans always remain in the loop:** a single button can stop one robot or the entire site.

It is like the Marvel universe, where Tony Stark tells J.A.R.V.I.S. what he wants, and the system turns his intent into a sequence of actions. The endpoint of construction technology's evolution looks similar, except that the result will be a building rather than a suit.

7. Why SIGNAX?

Today, SIGNAX is developing as a **Construction AI Harness**: data, knowledge, processes, BIM, AI tools and access rights are brought together in one environment, with AI agents working within it rather than being added as a separate layer on top of an existing CDE. As the focus shifts from digital processes to coordination of the physical state of the asset, this architecture will provide the foundation for a Construction Operating System.

7.1 Construction AI Harness Modules

Module	Harness Function	Today	In Development
S HUB	Infrastructure and security	Authentication, role-based access, event logs, backup and recovery, metadata	AI agent permissions and control
S DOCS	Data	Creation, storage and exchange of BIM models, drawings, specifications, cost estimates, schedules and contracts; all major formats	Automated document processing pipeline
S FORMS	Processes	Configurable planning, review, approval, work certification and reporting workflows: sequential, parallel and conditional	Machine-readable processes managed by agents

S BIM	Design and space	Proprietary 3D engine; viewing large models on a smartphone, clash detection and data export; attribute editing, custom parameters and documents linked to elements, such as site instructions and acceptance certificates. Simple geometry operations, such as splitting structures into work zones to record partial completion.	Generative design; spatial context for robots
S PROJECT MANAGEMENT	Operations	Production planning, task assignment and monitoring	Entity state history over time, linked to the schedule
S INSPECTION	Verification	Inspection and confirmation of completed work	Monitoring site events, camera and sensor data
S DASHBOARD	Analytics	Data collection across all modules, reporting and analytics	Predictive risk analytics
S AI Agent	Reasoning	Finding errors in documentation, preliminary review of design solutions, anomaly detection and schedule generation from design data	Risk assessment, agentic process management and management decision support
S KNOWLEDGE	Knowledge	File metadata: name, discipline, author, version, date and status	Vectorization and indexing on SIGNAX's own servers; project, company and external context memory
Integrations	Integrations	—	ERP, cost estimating and scheduling systems, IoT and robotic equipment

7.2 Data and Sovereignty

Data ownership. The data belongs to the client company. In the cloud (SaaS) version, SIGNAX does not use it without the client's consent. In the on-premises deployment, data remains within the client's own environment: SIGNAX cannot see or access it either.

Training models on client data. One client's data is not used to improve models for another. Anonymized data may be used to train SIGNAX's own models only under a separate agreement with the specific client.

Boundaries between clients. The SaaS version uses logical tenant isolation in a multitenant architecture: the backend strictly defines which users can access which companies and which operations they can perform, and checks permissions on every request. In the on-premises version, each client has its own database and file storage.

What leaves the system.

- **Model viewing.** The proprietary S BIM 3D engine processes models within SIGNAX. If a client connects Autodesk Forge Viewer, models are sent to Autodesk's servers for conversion.
- **External language models.** When OpenAI models are used for documentation review and assistant functions, the transmitted data is processed by the provider under the terms of its API policy.
- **Vectorization and indexing** are performed on SIGNAX's own servers using open-weight models in the 30B+ parameter class, such as the Qwen family.
- **Agent work patterns** are stored within the system as skills: Markdown descriptions with examples and sample output files.

Deployment. On-premises deployment serves companies with strict information security requirements; cloud deployment serves those prioritizing ease of rollout and support. SIGNAX operates in global clouds as well as local cloud environments in the UAE, Singapore, Malaysia, Vietnam and Kazakhstan. As companies pay increasing attention to data sovereignty, this architecture forms a **sovereign AI-ready Construction AI Harness**: an environment that will evolve into a **Construction Operating System** over time.

7.3 Key Differentiators

1. **Harness architecture** instead of a standalone AI assistant. AI gets not only access to documents, but also a managed context, project memory, specialized tools, machine-readable processes and structured data. The LLM is used for understanding, reasoning and orchestration, while platform algorithms perform deterministic work.
2. **Critical infrastructure at the core.** Security, reliability, authentication policies, access control and data sovereignty are built into the platform's foundation rather than added later.

3. **A proprietary 3D engine and SG format.** Large BIM models can be opened on a smartphone; geometry and properties are stored in the AI-accessible SGBIM + SGPROPS format.
4. **Structured entities and legal validity.** Documents and processes are defined by a set of attributes; e-signatures, "Release for Construction" stamps, version QR codes, and an immutable audit trail ensure legally binding document workflows today — and serve as a bridge during the transition until processes become fully digital.
5. **An AI agent with measured results.** The built-in agent already checks documentation, finds errors and generates certain types of documents from others, for example, schedules from drawings, which helps organize processes.
6. **A transparent data policy.** Data belongs to the client, is kept separate between clients and is not used for training without specific consent.

7.4 Roadmap

- **Today**: we are shaping the foundational Construction AI Harness. An AI agent reviews documentation, finds errors, works with structured data and generates new documents.
- **Tomorrow**: SIGNAX becomes a shared working environment for humans and AI. Agents employ project memory, specialized tools and machine-readable workflows to help humans set tasks and manage processes.
- **The day after tomorrow**: as the platform transitions toward direct coordination of physical construction operations, a Construction Operating System emerges. Through it, AI manages the creation of a physical asset under human supervision, coordinating many specialized agents, people and automated equipment while comparing the digital project representation with the actual state of the asset.

07



FIG. 10. **SIGNAX roadmap: today — tomorrow — the day after**

Today: the foundational Construction AI Harness. Tomorrow: a shared working environment for humans and AI. The day after tomorrow: a Construction Operating System coordinating physical operations.

Conclusion

Construction software was originally built primarily for humans and information management.

The next generation of systems must be designed for humans and AI together:

- **CDEs organized information storage and exchange;**
- **workflows connected information with project participants;**
- **the Construction AI Harness brings data, knowledge, processes and software tools into one environment in which humans and AI work together;**
- **the Construction Operating System will understand the state of the project and coordinate what needs to happen next.**

In such an architecture, the LLM becomes the intelligence layer that understands the task, reasons and orchestrates actions. Specialized tools perform precise queries, calculations, checks, format conversions and other deterministic operations.

The Construction AI Harness determines which context the model receives, which tools it can access, which actions require human confirmation, and how results are returned to project memory.

We believe this architecture will define the next stage of CDE development.

As the environment around the LLM—the **Construction AI Harness**—receives increasing amounts of information about the physical state of the asset and coordinates digital processes as well as the actions of humans, equipment, autonomous agents and robots, it will evolve into a **Construction Operating System**: an environment for managing physical construction operations.

SIGNAX is being built as this kind of platform from the outset: a new category of digital infrastructure that brings project data, project and company memory, processes and

tasks—knowledge about construction production—together with artificial intelligence in an AI-native environment. The establishment of this infrastructure aims to facilitate the industry's transition from digital document management to AI-managed construction operations.

Appendix A. Conducting an Experiment in Reviewing Design Documentation Using AI Models

In September 2026, SIGNAX engineers conducted a series of experiments reviewing the same set of design documentation.

Subject: "Structural Design" section (40 PDF drawings) for a manufacturing building with an area of 21,600 m².

Method: the design documentation set was reviewed by different LLMs using different AI workflows. A total of seven runs were performed—ranging from complex, multi-stage analysis lasting several hours to fast automated reviews using modern multimodal models.

A reference set of 14 significant errors was selected for comparison. This was not a full list of errors in the documentation, but covered the most obvious, typical cases:

- a unique item number is used for different elements;
- identically designated parts have different specifications;
- drawing data conflicts with the bill of materials;
- template placeholders remain unfilled;
- some dimensions are incomplete;
- a reference points to the wrong sheet or assembly;
- an arithmetic expression contains an error;
- a component's mass does not match its geometry or specification;
- element quantities differ between tables;
- a value on one sheet contradicts the summary schedule.

Criteria: the focus was on how useful the result was to an engineer, rather than how many findings a neural network could generate:

- how many findings could actually be confirmed;
- how many significant problems AI missed;
- whether AI could connect information across different sheets;
- whether it could check arithmetic, specifications and cross-references;
- whether it could distinguish a clear error from ambiguity;
- how much manual review work was reduced.

The results showed that review quality is determined by more than just the model. The process of analyzing the documentation is equally important.

Best-performing model: GPT-6 Astra delivered the best results in the experiment. It generated 23 findings: an engineer confirmed 21, one was a false positive, and one was classified as debatable.

The model also detected 13 of the 14 reference defects and identified several problems absent from previous reports.

Thus, the final results of the experiment are as follows:

Metric	Value
Total issues found	23
Confirmed by an engineer	21
False positives	1
Debatable	1
Precision	21 / 23 = 91.3%
Completeness by the reference set	13 / 14 = 92.9%
Manual engineer review	220 min
Model review	10 min

Engineer validation of model results	25 min
Reduction in total process time	185 min, or 84%
Engineer time savings	195 min, or 88.6%

Conclusions.

1. **Documentation review** is a system-level task, rather than the work of a single neural network. Quality depends on both the model and the analysis process.
2. **The number of findings says nothing about quality**. The weakest model, Qwen3-VL 8B, produced 47 findings, almost all false. The strongest produced 23, of which 21 were confirmed, and found 13 of the 14 reference errors.

Appendix B. Glossary

Term	Meaning
LLM (Large Language Model)	A neural-network-based artificial intelligence system designed to deeply understand, process and generate natural-language text.
AI-native	Data and systems designed from the outset for human–AI collaboration: structured, machine-readable, decomposed into parts and verifiable
CDE	Common Data Environment: a shared project data space for all participants
Construction AI Harness / Harness	An AI-native software environment around an LLM that organizes data and memory, manages context, provides specialized tools and machine-readable workflows, performs deterministic operations and controls AI actions. The model within the Harness reasons and orchestrates; algorithms perform precise operations; people make critical decisions.
Construction Operating System / COS	A stage in the development of the Construction AI Harness at which the system coordinates digital data and processes as well as the physical state of the asset and construction operations: the work of people, agents, equipment and, eventually, robots.
MCP	Model Context Protocol: an open protocol for connecting AI models to external tools and applications
SDD	Spec-driven development: an approach in which a person defines the specification, an agent executes it, and the person accepts the result
SG format	SIGNAX's family of AI-native formats: SGBIM for geometry, SGPROPS for BIM properties, SGMPP for schedules, and others
Vector database	A store in which document fragments are represented as vectors for semantic search
Context	The information available to the model while performing a task; a limited resource. A key function of the Harness is to select the minimum sufficient context rather than pass the entire project dataset to the model.

Project / company / external context memory	Knowledge base levels: data from a specific project, accumulated company experience, and external regulations and reference sources
NCR	Non-Conformance Report: a notice of nonconformity
RFI	Request for Information: a request for clarification or additional information
Git	A distributed version control system.

About the Company

SIGNAX Software Design LLC provides a digital construction management platform. The company was founded by specialists in construction management and supervision.

SIGNAX is based in Dubai, UAE, and Singapore, and works with clients, contractors, consultants and engineering companies on projects ranging from commercial and residential real estate to major infrastructure and energy facilities across the Middle East, Southeast Asia and the CIS.

signax.io
info@signax.io
+971 50 138 0516

© 2026 SIGNAX Software Design LLC. All rights reserved. This document is provided for informational purposes and reflects the current direction of SIGNAX product development. Forward-looking statements are subject to change. All trademarks mentioned are the property of their respective owners.