Towards Practical Digital Twin Construction Systems

Timson Yeung[1] and Rafael Sacks[1]

[1]Seskin Virtual Construction Laboratory, Technion – Israel Institute of Technology, Israel
timsonyeung@campus.technion.ac.il, cvsacks@technion.ac.il

Abstract

Building construction projects are socio-technical production systems in which variance, uncertainty and unpredictable events are the norm. Such systems need short-cycle Plan–Do–Check–Act (PDCA) loops to detect anomalies in production flow as early as possible so that remedial actions can be evaluated and well-informed decisions made. This is the goal of Digital Twin Construction (DTC). We define practical DTC as a minimal coherent system that captures a project's current status, compares the status to design and planning intent, analyses variance, predicts the outcomes of possible interventions, and thus informs specific managerial decisions at the required cadence. We introduce a 'thinking backward' decision-centric design approach for DTC systems that traces requirements from a target decision and its PDCA cycle through knowledge and information needs to data acquisition and enabling technologies. The approach characterizes managerial decisions by autonomy, fidelity, and scope, and proposes five degrees of implementation as system capability configurations for incremental development. This contributes an actionable basis for specifying and incrementally implementing DTC systems that deliver decision value early while maintaining a clear pathway toward more advanced implementations.

Keywords

Digital Twin Construction; Automation; Lean Construction; Decision-centric System Design; Plan-Do-Check-Act; Knowledge-Information-Data

1 Introduction

Building construction projects are socio-technical production systems where variance, uncertainty, and unpredictable events are the norm [1]. Production rates fluctuate, defects trigger rework, and small human errors can compound into major delays, while design changes, technical constraints, and supply-chain disruptions add further uncertainty. Managing such conditions requires short-cycle PDCA loops [2] to detect deviations early and support timely corrective decisions while recovery remains possible.

DTC [3] addresses this need. It is a data-centric mode of construction management that combines digital-twin information systems with lean construction thinking to close production control loops and support PDCA during the construction phase. DTC workflow continuously compares project intent with updated project status and feeds the resulting evaluations to people or algorithms that intervene on site and verify outcomes in the next cycle.

State-of-the-art research shows progress toward realizing DTC workflows, including interoperable information management layers, increasingly automated progress monitoring. Case studies have demonstrated closure of intent to status to intent feedback loops in live projects [4,5]. These efforts reveal that the key challenge in implementing DTC systems is not the availability of sensing, analytics, or cloud infrastructure, but the design problem of turning these components into coherent decision support pipelines that reliably improve planning and control outcomes. Whereas successful case studies begin from a clearly bounded use case and demonstrate measurable benefits, many other initiatives struggle to translate technological capability into repeatable closed-loop decision making.

A common pattern in both research and practice is a predominantly technology-driven starting point: teams begin from available technologies and only later ask what decisions the system should improve, so that digitalization often fails to deliver full decision value.

We propose a decision-centric perspective for designing practical DTC systems that support incremental adoption without sacrificing coherence. This is a "thinking backward" approach [6] that traces requirements from a target decision and its PDCA cycle back through knowledge and information needs to data acquisition and enabling technologies. It characterizes production planning and control decisions along three dimensions, autonomy, fidelity, and scope, and presents degrees of DTC implementation as coherent system configurations that help teams set realistic targets, avoid overengineering, and plan a stepwise pathway toward stronger closed-loop control.

2 State-of-the-Art in Digital Twin Construction

Sacks et al. [3] introduced DTC as a data-centric mode of construction management that combines digital-twin information systems with lean construction thinking to close production control loops and support PDCA cycles during the construction phase. In the DTC workflow, design and planning intent is formalized as project intent information (the as-designed product and as-planned process). Project intent is continuously compared with project status information (the as-built product and as-performed process), which is derived from raw data from site and supply-chain monitoring that is interpreted using data-fusion functions. Subsequently, the comparisons yield conformance evaluations that inform project status knowledge suitable for proposing possible corrective interventions. In parallel, simulations and analytics produce new project intent knowledge that supports performance prediction and the evaluation of decision alternatives, while bidirectional feedback connects these digital representations back to both planning/design updates and real-time production control on site. Archiving digital building twins that capture all this information enables learning across projects.

Practical DTC systems require a system architecture with explicit data layers and interfaces that integrate sensing, management, and analytics into a decision support pipeline. A common approach uses Semantic Web technologies to link heterogeneous product, process, and status data through open standards. Schlenger et al. [7] addressed the limited guidance on data management layers for DTC by proposing a plug-in-based reference architecture that couples multiple RDF graphs with specialized databases and introduces a Digital Twin Construction Ontology to compare planning information with monitored status, validated on the ConSLAM dataset. Complementing this platform view, Zheng et al. [8] proposed the DiCon-DT framework for construction workflow monitoring, where an ontology network integrates diverse data into a semantic DT data lake and supports simulation and contextual interpretation.

Among the technical building blocks of DTC systems, automated progress monitoring has been developed rapidly. Recent studies demonstrate that imagery and video can support inference of construction progress [9]. For geometry centric assessment, point cloud pipelines continue to improve scan to BIM alignment and association for element level completion inference [10]. Complementary nonvisual sensing provides additional evidence streams that can help connect captured evidence to coordination and compliance actions in management workflows [11].

Beyond monitoring, researchers have connected digital twin status synchronization with simulation and optimization to enable decision models to update assumptions and explore alternatives using current constraints. Work in this area typically targets two horizons: production planning that supports system configuration and medium-term plans, and production control that adjusts daily logistics and execution in near real time. For example, Yeung et al. [12] integrated an ontology based digital twin with parametric agent based simulation to support adaptive production system design, enabling real time data ingestion, status model instantiation, and scenario analysis. Jiang et al. [13] proposed a digital twin-enabled real time synchronization system for just-in-time precast on site assembly that uses real time resource status and construction progress to support planning, scheduling, and execution coordination and dynamic adjustment.

Large scale research and innovation initiatives have also sought to integrate the DTC building blocks into end-to-end prototypes and to validate them beyond isolated case studies. Under the EU Horizon 2020 Digital Building Twins call, a coordinated cluster of projects moved construction phase digital twin concepts toward cloud prototypes that combine data integration with decision support services. BIM2TWIN aligned closely with the DTC system concept by specifying a Digital Building Twin platform based on semantic-linked data from multiple sensor streams into a real time project status model along with data from BIM and project management data, and exposes it through APIs and dashboards for an extensible set of construction management applications. ASHVIN and COGITO complement this platform logic with toolboxes and vertical services for monitoring, planning, and governed workflows, including quality and health and safety functions and traceability mechanisms such as smart contracts. BIMprove focused on real time tracking and field interfaces such as AR and VR to keep BIM and digital twin views synchronized during execution.

Beyond component-level advances, recent full-workflow demonstrators illustrate how DTC concepts translate into closed-loop production control in live projects. Valinejadshoubi et al. [4] implemented a closed-loop production control workflow for a cement silo project by integrating BIM, IoT sensing, and GIS to enable near-real-time progress monitoring, planned versus actual comparisons, and deviation-triggered management actions. Their system coupled a BIM model with an ArcGIS-based digital twin environment and a cloud database, and it streamed formwork sensing data to support timely conformance checking and corrective intervention. They reported a 28% reduction in execution time and a 15% reduction in construction wastage, which they attributed to faster information updates and tighter coordination enabled by the integrated platform.

Martinez et al. [5] presented a mechanized tunnelling demonstrator in which the digital twin closed an intent-to-status-to-intent loop for bracketry drilling design and installation. The team consolidated TBM and segment-manufacturing data into an as-built tunnel lining model and then used the updated status model to define no-drill zones, perform clash detection, and support bracketry design and execution. Automating this consolidation reduced the per-ring modelling effort from 15 to 20 minutes to about 1 minute of manual quality checks, cutting cycle time by more than 90%. The authors also noted persistent barriers such as imperfect interoperability across proprietary software and hardware systems and the lack of standardized BIM-IFC representations of tunnel geometry and information.

3 Decision-Centric System Design

The state-of-the-art review shows rapid progress toward realizing DTC workflows, including interoperable information management layers, increasingly automated progress monitoring, and end-to-end demonstrators that close intent to status to intent feedback loops in live projects. These efforts also reveal that the hardest part of DTC system implementation is often not the availability of sensing, analytics, or cloud infrastructure, but the design problem of turning these ingredients into a coherent decision support pipeline that reliably improves planning and control outcomes. Whereas successful case studies begin from a clearly bounded use case and demonstrate measurable benefits, many other initiatives struggle to translate technological capability into repeatable closed loop decision making.

A common pattern in both research and practice is a predominantly technology-driven starting point. Teams adopt tools such as cameras, LiDAR sensors, or monitoring services first and only later ask what decisions they improve and how they strengthen a closed-loop PDCA cycle. This opportunistic mindset is understandable. In research, it is often easier to trial novel sensing, algorithms, or system architectures than to define and validate end value across heterogeneous stakeholders. In commercial settings, technology vendors naturally lead with their product propositions and attempt to persuade project teams to adopt their products. As a result, many DTC implementations start from available technologies and then work forward toward decisions and learning loops.

We propose to flip this logic and take a decision-centred approach to DTC system design. DTC is framed as a management paradigm whose core purpose is to support and improve decisions through systematic learning, rather than as an umbrella label for digitalization technologies. The starting point is therefore the set of managerial decisions that are intended to be made within a closed loop PDCA cycle.

Two broad horizons are especially relevant in construction. Production control decisions govern near term execution, including daily coordination, constraint removal, sequencing, workface logistics, and the allocation of crews, equipment, and materials. Production planning decisions shape the production system itself, including adaptive production system design and the periodic reconfiguration of plans, control policies, and resource strategies in response to evolving constraints and performance. In both horizons, the central question is not what the digital twin can technically compute, but what decisions must be made, who makes them, how often they must be revisited, and what learning is expected from repeated cycles.

From an explicit target decision, the system designer can work down through levels of knowledge, information and data. At the knowledge level, a designer asks what knowledge about project intent and project status is needed to make and improve the decision over time, including what must be understood about conformance, deviations, causes, and consequences. At the information level, they ask what information must be synthesized to support those knowledge claims. For example, which state variables, constraints, progress states, resource states, and performance indicators must be explicitly represented. At the data level, they ask what combinations of raw observations can be used to derive that information. Multiple data sets can often support the same information requirements, and direct measurement is not always feasible. In such cases, indirect measurements, partial measurements, or complementary measurements from different tools may be combined through interpretation and data fusion to reconstruct the information of interest. Only once these requirements have been defined explicitly does it become rational to select monitoring solutions, data management strategies, and system architectures that implement a pipeline from data to information and then to knowledge that supports the decision defined at the outset.

This sequence helps avoid over engineering and over complication, because it focuses attention on the minimal sufficient evidence and transformations needed for the decision, rather than on maximizing the volume of captured data or assembling an attractive architecture that lacks a clear operational purpose.

A practical consequence of starting from decisions is that the use case must be characterized in a way that sets realistic expectations for the required system capability and complexity. As illustrated in Figure 1, production planning or control decisions can be described using three parameters with values arrayed along three axes: *autonomy*, *fidelity*, and *scope*. The autonomy axis describes the prospective agency of the human planner in the decision process, ranging from *human decide*, through *human decide with machine advice*, to *machine decide with human supervision*, and ultimately to *machine decide*, while recognizing that responsibility, risk, and centralized versus distributed authority create intermediate nuances. The fidelity axis describes how close the information must be to reality for the decision to be valid. Fidelity can be specified through required *precision* or *resolution*, required *accuracy*, *acceptable latency* and *update frequency*. The scope axis describes the breadth and interdependence of the decision within the production system, including whether it is *localized or global*, whether it is a *single decision or a coupled set of decisions*, and whether it creates dependencies and cascading effects across other decisions and system states. Together, these axes help determine whether a simple single component solution is sufficient or whether the use case demands multiple modules, rich data fusion, and/or sophisticated predictive or prescriptive analytics.

We can apply this decision-centric approach to the silo case study [4] reviewed in the previous section. In slip-form construction, the team can first isolate the near-term production-control decisions that plausibly benefit from short-cycle PDCA. The silo workflow supports four such decisions: 1) whether to adjust formwork alignment to maintain verticality and rotation, 2) whether to intervene when curing conditions drift toward risk, 3) whether to change pouring rate and near-term resource allocation to maintain planned progress, and 4) what materials or embedded components to pull to the site next as progress advances. As a group, these decisions remain in the *human decide* autonomy class. They stay within a small and localized scope centred on the silo production-control loop. The use case demands high fidelity, with minute-level temporal resolution and high accuracy, and an hourly update frequency that matches the production-control cadence.

With the decision set fixed, the designer can then trace down the knowledge, information, and data pyramid to elicit system requirements. Figure 2 summarises this trace as a dependency map. At the knowledge level, the four decisions require knowledge of structural alignment relative to intent, curing condition status relative to thresholds, progress versus plan, geometric conformance for embedded elements, and material pull needs for the next work segment. At the information level, these knowledge requirements translate into a minimal set of information products that the system must maintain, including real-time construction parameters, trend histories that expose drift, visual status indicators on the model, progress metrics, alerts that flag excursions, and an as-built versus as-designed overlay that supports conformance assessment. At the data level, the same trace identifies a minimal viable data set and its sources, including distance and inclination sensing for geometry and stability, temperature sensing for curing conditions, BIM intent geometry as the reference for planned states, timestamps to support trend and auditability, and site spatial context where needed for situational framing.

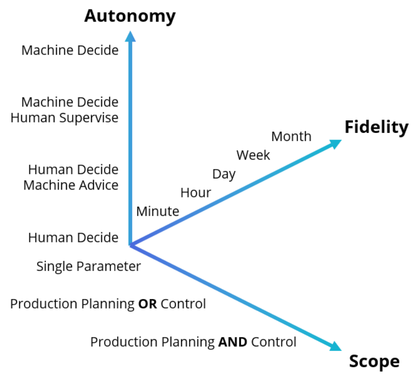

Figure 1. Orthogonal axes characterizing a DTC system use case

This worked example shows how the thinking-backwards approach helps convert a case description into an explicit requirement specification for a practical DTC system. By linking a target decision and its cadence to concrete knowledge, information products, and data sources, the approach guides purposeful technology selection and, equally, prepares the production team to use the resulting information routinely to check conformance and act within the decision window.

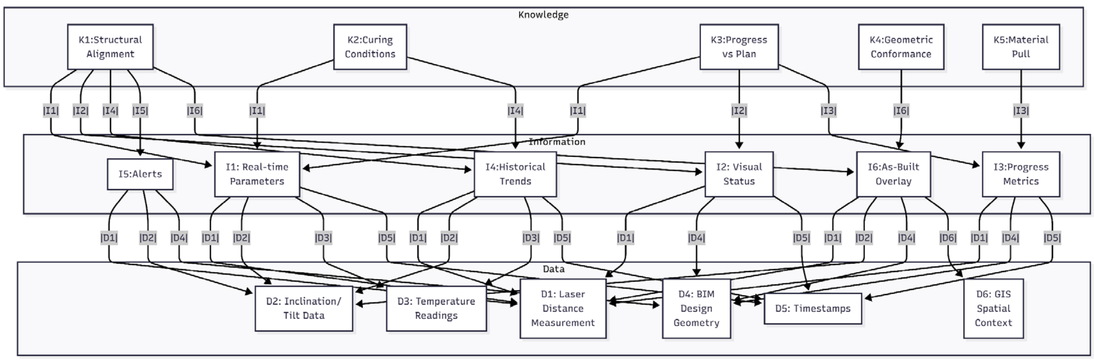

Figure 2. Derivation of required knowledge, information, and data for use cases in Valinejadshoubi et al. [20]

4 Degrees of DTC System Implementation

Building on the decision-centric framing, a key practical question is, "Given a target decision, its requirement for autonomy, fidelity, scope, knowledge, information, and data, what degree of DTC implementation is realistic, and what capabilities must be in place for the PDCA loop to run systematically?" The supportive and autonomous visions of DTC provide useful long-term reference points, but they can also appear overly ambitious when viewed against the realities of construction variability, fragmented responsibilities, and the cost of instrumentation and integration. At the same time, meaningful DTC is not an all-or-nothing proposition. Projects can achieve measurable improvement with partial digitization and partial automation, provided that the system delivers decision support that is traceable to the use case and enables learning across repeated cycles.

This section therefore proposes five degrees of DTC implementation (Figure 3) that describe coherent capability bundles, rather than a checklist of isolated technical features. The degrees presented in Figure 3 are not intended as a strictly linear maturity ladder. Organizations can adopt different degrees for different decisions, and a single project can exhibit mixed degrees across planning horizons and work packages. What changes across degrees is primarily the extent to which the system implements the knowledge-information-data chain for the target decision, and the extent to which monitoring, interpretation, and feedback become structured, repeatable, and timely. Higher degrees typically correspond to higher attainable autonomy, higher fidelity of the intent and status representations, and broader scope of the decision support.

The first degree is **unassisted**. Here, the PDCA loop exists in the traditional sense described in lean construction and production control practice, but it is largely manual and relies on managerial judgement. The only structured digital representation of intent is usually a schedule and associated planning artifacts, while project status awareness is formed through site walks, phone calls, ad hoc updates, and meetings. Teams check progress against the plan and revise near-term commitments, but the evidence base is fragmented and informal, so the loop's effectiveness depends on the availability, memory, and interpretation capacity of the management team.

The second degree is the **digital model**. Relative to the unassisted case, teams deploy monitoring technologies and generate richer project status data, for example through photogrammetry and computer vision, scanning, tagging, or other capture methods that observe both inputs and outputs of the production system. These observations can indicate where labour and materials have been deployed and what intermediate or final products have been installed in specific work zones. The key characteristic is that the data typically flows into the human planner's situational awareness model rather than into a structured digital interpretation pipeline. Project intent information still resides primarily in the schedule, BIM, and planning documents, and the comparison of intent and status remains a human task. This degree is therefore best described as data-augmented rather than systematically data-driven. Richer, more frequent, and potentially more objective observations support decisions, but they do not yet consistently produce consolidated project status information and actionable knowledge.

The third degree is **digital shadow**. This degree retains monitoring of inputs and outputs, but it differs from the digital model in that the system integrates multiple data sources and performs explicit interpretation and data fusion to generate consolidated project status information. Different sensing modalities provide partial views of the production system, and their combination makes it possible to reconstruct a more complete picture across resources, processes, and products. For example, beacons can provide evidence of labour presence and

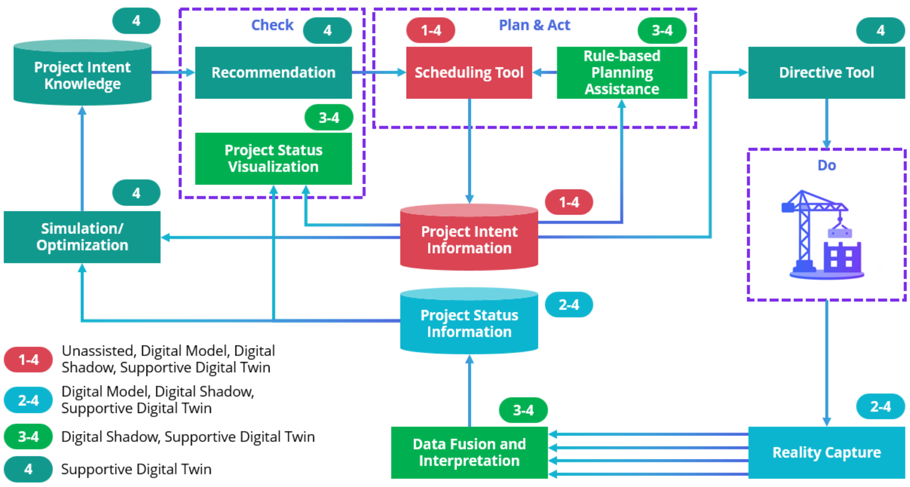

Figure 3. Degrees of DTC system implementation

movement, RFID can indicate material location and consumption, image-based methods can capture visible installation progress, and aerial or terrestrial scanning can characterize geometry for exterior or structural elements. The system uses these complementary observations to support consistent conformance evaluation, not only identifying deviations between intent and status but also helping explain why deviations occur, such as delays driven by missing inputs, rework, or disrupted workflows. Importantly, the feedback loop is increasingly explicit. Updated project intent information, including revised plans and constraints, is recorded and used as input for subsequent evaluations, enabling rule-based planning assistance and systematic iteration across PDCA cycles.

The fourth degree is the **supportive digital twin**. At this degree, the system extends beyond consolidated status to generate project intent knowledge through simulation, analytics, and optimization. The digital shadow foundation provides the evidence base for current state estimation, while predictive and prescriptive modules explore what-if scenarios, estimate consequences, and propose decision alternatives. This corresponds to predictive situational awareness, where the system supports not only understanding what happened and why, but also anticipating what will happen next and evaluating the likely impact of alternative actions. The system remains supportive because the human planner retains decision authority. It provides recommendations, explanations, and visualizations that help the planner decide whether and how to act, and the resulting plan updates are fed back into digital representations to support continuous improvement.

The fifth degree is the **autonomous digital twin**, which can be described as a digital agent governing decision selection. Here, the fused project status information and the project intent information feed an autonomous decision algorithm that evaluates alternatives and selects an action, which is then written back into the project intent representation and issued as directives, policies, or control inputs for execution. In the ideal form, the planner no longer participates in the loop. However, in practice, construction conditions, risk, and capital intensity will motivate a supervised autonomy arrangement in which human planners oversee, approve, or override algorithmic decisions. This makes explainability, traceability, and logging critical, both for operational assurance and for long-term learning through a digital repository of intent, status, and decision histories that can improve future models and algorithms.

Across these degrees, the central point remains consistent with the preceding section: the PDCA loop is not binary. Lower degrees can still run PDCA cycles, but with higher reliance on intuition and informal evidence. Higher degrees make the loop more systematic by formalizing what must be known, what must be inferred, how intent and status are compared, and how actions are selected and recorded. This typology therefore provides a practical bridge from a decision-centric system design approach to implementable system targets.

In relation to existing digital twin maturity models, these degrees align broadly with established capability progressions in the literature, particularly the transition from digital model and digital shadow toward predictive or prescriptive and autonomous twins described in models such as [14] and [15]. The similarity lies in the increasing formalization of data integration, interpretation, and feedback, and in the resulting shift from status visibility toward prediction, decision support, and eventually decision automation. The novelty of the proposed degrees is that they recast this progression for construction planning and control around the implementation of a target PDCA loop rather than around digital twin technological maturity. This shift is reflected first in the inclusion of an explicit unassisted degree, which recognizes that PDCA already exists in practice before a formal digital twin is introduced and therefore starts the framework from the real managerial baseline rather than from the first digital artifact. It is reflected again in the definition of the higher degrees, which are distinguished not simply by more advanced technology but by how systematically the knowledge-information-data chain is implemented to support a specific managerial decision.

5 Discussion

DTC initiatives frequently begin with a technology-first mindset, where teams adopt sensing and analytics tools and then attempt to fit management decisions around whatever data those tools produce. This approach often results in fragmented architectures and unclear value, because it optimizes technical capability rather than strengthening a specific PDCA loop. The framework proposed in this paper reverses that sequence. It starts from a decision that must be improved, and it treats system components as means to close that decision loop reliably.

The autonomy, fidelity, and scope characterization helps developers translate decisions into implementable requirements. Autonomy clarifies who decides and what the system is allowed to do, fidelity clarifies how accurate and timely the information must be, and scope clarifies how far the decision's effects propagate across planning and control. Framed this way, a project team can state what the system must deliver before selecting technologies, which reduces over-engineering and helps align organisational responsibilities, data ownership, and update frequency with the actual decision context.

Once requirements are clear, the knowledge–

information–data decomposition provides a practical route from decision needs to system design. It derives the minimal data capture and processing pipeline from what the decision maker must know about project intent and project status, rather than collecting data opportunistically. This logic supports lean implementation, including the use of partial measurements, proxies, and staged sensing deployments, provided they satisfy the fidelity requirements of the targeted PDCA loop.

Finally, the degrees of DTC implementation position DTC as incremental progress rather than a yes or no achievement. They enable teams to diagnose what a system can currently support and to prioritize the next feasible capability that materially strengthens the decision loop, without assuming that full automation or comprehensive sensing is necessary. The remainder of this discussion illustrates this practicality by applying the framework to three real-world demonstrators.

Table 1 details the proposed characterization with respect to three case studies and shows, first, how each implementation starts from a clearly bounded managerial decision and, second, how the system architecture is designed to deliver the specific status information needed to act on that decision.

Valinejadshoubi et al. [4] targeted concreting crew task assignment. While decisions remained human-centric (human decide), the use case demanded high fidelity with minute-level resolution. Consequently, the system required automated fusion of formwork sensing data to trigger threshold-based recommendations, placing it at the Digital Shadow level (Degree 3).

Similarly, the mechanized tunnelling system [5] supported bracketry design. It operated at a lower temporal cadence (daily) but required precise spatial data fusion. By consolidating TBM and manufacturing data into an as-built model to identify no-drill zones, it provided machine advice, also qualifying as a Digital Shadow.

As part of the EU BIM2TWIN research project, the authors implemented a DTC system for adaptive production system design rather than short-cycle production control [12]. The system integrated an ontology-based digital twin with agent-based simulation to ingest current project status, instantiate the status model, and evaluate alternative production configurations under changing project conditions. Relative to the first two case studies, it added predictive recommendation and supported updating the baseline production plan, but it did not rely on automated reality capture and only partially automated conformance evaluation. It therefore sits between Degree 3 and Degree 4, showing that the proposed degrees are better understood as practical capability regions along a spectrum than as rigid thresholds that a project must fully satisfy before it becomes useful.

Table 1. Classification of use case and degree of DTC implementation for three case studies

Study		Valinejadshoubi et al. [4]	Martinez et al. [5]	Yeung et al. [12]
Use Case		Concreting crew task assignment, conformance checking	Define no-drill zones, perform clash detection, and support bracketry design and execution	Adaptive production system design through adjusting resource, location, and work package parameters
Autonomy		Human decide	Human decide with machine advice	Human decide with machine advice
Fidelity	**Temporal Resolution**	Minute	Day	Day
	Update Frequency	Hourly	Daily	Weekly
	Accuracy	High	High	Medium
Scope		Small	Medium	Medium
Reality Capture		Yes, automated	Yes, automated	No
Data Fusion and Interpretation		Yes, automated	Yes, automated	Yes, automated
Conformance Evaluation		Yes, automated	Yes, automated	Yes, semi-automated
Simulation/Optimization		No	No	Yes, agent-based simulation
Recommendation		Yes, remedial action if monitored concrete condition below preset threshold	Yes, no drill zones per segment	Yes, alternative production system configurations
Report & Visualization		Yes, sensor data visualization and concrete pouring progress	Yes, 3D as-built IFC model with production information per segment	Yes, schedule, progress, and simulation results of alternative plans in visual dashboard
Rule-based Assistance		No	No	No
Directive Tool		No	No	Yes, update new baseline production plan to digital twin repository
Degree of DTC Implementation		3 - Digital Shadow	3 - Digital Shadow	Between 3 – Digital Shadow and 4 – Supportive Digital Twin

Taken together, the case studies show that practical DTC implementation is already feasible with existing technologies when the system is designed around a bounded managerial decision and its associated PDCA loop. The reported outcomes show that meaningful operational value can be achieved without an all-encompassing or highly autonomous twin, provided that the implementation delivers the specific status information, interpretation, and decision support required for that use case with sufficient reliability and cycle time. At the same time, the framework is most readily applied when the target decision and its PDCA loop can be defined clearly. In more highly coupled project settings, applying it may require decomposing interdependent decisions into smaller linked loops before coherent system requirements can be specified. Seen this way, the proposed degrees are not targets to be pursued for their own sake, but a framework for characterizing which capabilities have been implemented, which remain absent, and which additional capabilities are justified by the expected decision value, thus reframing DTC from an idealized end state to a practical system design problem of aligning scope, automation, and decision support with concrete managerial needs.

6 Conclusion

Short-cycle PDCA is needed to manage variance and uncertainty during construction production, yet DTC systems are often conceived based on available technologies and idealized end states rather than being based on the decisions they must support. This paper proposed a decision-centric design perspective and the notion of practical DTC as the minimal system capability required to improve a specific PDCA decision at its required cadence.

The paper introduced three decision axes, derived design principles that trace requirements from the target decision through information and knowledge needs to data acquisition and defined five degrees of DTC implementation to guide staged development. Applied to three real-world case studies, the framework separated the managerial use cases from the particular technologies and architectures used to implement them and clarified that PDCA improvement is achievable at multiple degrees of digitalization and automation. Together, these contributions provide a pragmatic basis for specifying and incrementally implementing DTC systems that prioritize decision value over technology completeness.

References

[1] M. Abdelmegid, A. Tezel, C.O. Sandoval, Z. Fang, W. Collinge, Conceptualizing Digital Twins in Construction Projects as Socio-Technical Systems, ISARC Proceedings 2024 Proceedings of the 41st ISARC, Lille, France (2024) 577–584.

[2] W.E. Deming, Out of the crisis, reissue, MIT press, 2018.

[3] R. Sacks, I. Brilakis, E. Pikas, H.S. Xie, M. Girolami, Construction with digital twin information systems, DCE 1 (2020) e14.

[4] M. Valinejadshoubi, R. Sacks, F. Valdivieso, C. Corneau-Gauvin, A. Kaptué, Digital Twin Construction in Practice: A Case Study of Closed-Loop Production Control Integrating BIM, GIS, and IoT Sensors, Journal of Construction Engineering and Management 151 (2025) 05025014.

[5] E. Martinez, F. Stauch, B. García De Soto, R. Sacks, Integration of building information modelling and digital twin in mechanized tunnelling construction, ECAM (2025).

[6] B. Flyvbjerg, D. Gardner, How big things get done: the surprising factors that determine the fate of every project, from home renovations to space exploration and everything in between, Crown Currency, 2023.

[7] J. Schlenger, K. Pluta, A. Mathew, T. Yeung, R. Sacks, A. Borrmann, Reference architecture and ontology framework for digital twin construction, Automation in Construction 174 (2025) 106111.

[8] Y. Zheng, A. Al Barazi, O. Seppänen, H. Abou-Ibrahim, C. Görsch, Semantic digital twin framework for monitoring construction workflows, Automation in Construction 176 (2025) 106301.

[9] F. Pfitzner, S. Hu, A. Braun, A. Borrmann, Y. Fang, Monitoring concrete pouring progress using knowledge graph-enhanced computer vision, Automation in Construction 174 (2025) 106117.

[10] Y. Zhang, L. Lu, X. Luo, J. Pan, Global BIM-point cloud registration and association for construction progress monitoring, Automation in Construction 168 (2024) 105796.

[11] M. Khazen, M. Nik-Bakht, O. Moselhi, Monitoring workers on indoor construction sites using data fusion of real-time worker's location, body orientation, and productivity state, Automation in Construction 160 (2024) 105327.

[12] T. Yeung, J.G. Martinez, J. Schlenger, A. Borrmann, R. Sacks, Integrating digital twin and agent-based simulation to support adaptive production system design in building projects, Automation in Construction 180 (2025) 106550.

[13] Y. Jiang, M. Li, M. Li, X. Liu, R.Y. Zhong, W. Pan, G.Q. Huang, Digital twin-enabled real-time synchronization for planning, scheduling, and execution in precast on-site assembly, Automation in Construction 141 (2022) 104397.

[14] A. Fuller, Z. Fan, C. Day, C. Barlow, Digital Twin: Enabling Technologies, Challenges and Open Research, IEEE Access 8 (2020) 108952–108971.

[15] W. Kritzinger, M. Karner, G. Traar, J. Henjes, W. Sihn, Digital Twin in manufacturing: A categorical literature review and classification, Ifac-PapersOnline 51 (2018) 1016–1022.