A THREE-LAYER TAXONOMY AND PREDICTIVE DIGITAL TWIN FOR ROLE-AWARE HUMAN–ROBOT COLLABORATION IN 3D CONCRETE PRINTING

Ibrahim Chahbaz[1], Reina Chmayssany[1], Jiany Samara[1], Noha Ramadan[1], and Hiam Khoury[1]
[1]American University of Beirut, Beirut, Lebanon

Abstract

3D concrete printing (3DCP) has reached early industrial deployment, yet current standards exclude environmental, health, and safety aspects of robotic operation, and existing collaborative-robot frameworks do not address process-coupled hazards or multi-role construction environments. This paper proposes a three-layer Human-Robot-Site taxonomy and a site-wide predictive digital twin with five-second update cadence and short-horizon forecasting. The framework is operationalized through seven role-aware interaction patterns aligned with ISO 10218. A scenario-based evaluation of the Human-layer physiological pipeline using a proxy dataset demonstrates the feasibility of stress classification and calibration. The work provides a reference architecture for human–robot–site integration in 3DCP.

Introduction

3D concrete printing (3DCP) has crossed an inflection point. Recently, Harcourt Technologies completed the first social-housing project compliant with ISO/ASTM 52939:2023, the new qualification standard for additive construction, using a gantry printer to deliver three terraced units in 132 working days; a 35% reduction relative to conventional methods (COBOD International, 2024). The platform is no longer experimental.

What has not crossed that inflection point is the safety framework around it. For instance, Section 1 of ISO/ASTM 52939:2023 is unambiguous: the standard *"does not cover environmental, health* and *safety aspects that apply to printing facility setup, material handling, operating of robotic equipment, and packing of equipment and/or elements for shipping"* (ISO/ASTM, 2023). The standard governing print quality says nothing about how to keep workers safe alongside the machine that produces it. The companion standards available to fill this gap (ISO 10218-1:2025 and ISO 10218-2:2025) integrate the former ISO/TS 15066 collaborative requirement and were written for stationary articulated industrial robots in fixed cells (ISO, 2025). They define collaborative applications around four operational types (e.g. monitored standstill, hand guiding, speed and separation monitoring, and power and force limiting) but presuppose a workspace topology and hazard model that 3DCP does not match.

3DCP differs from cell-based collaborative robotics in four ways that matter for human–robot collaboration (HRC). First, the dominant on-site platform is a gantry: large, slow, with high inertia, traversing an open construction site rather than operating within a guarded cell. Second, the robot is tightly coupled to a continuous material process (e.g. pump pressure, nozzle flow, fresh concrete rheology) whose state cannot be paused without quality consequences such as cold joints, where bond strength reduces substantially with delay (Panda et al., 2018). Third, multiple worker roles operate concurrently in or near the work envelope. In other words, print operators monitor the system from a console, rebar placers physically intervene between layers, material handlers tend the pump and supply, and supervisors and other construction personnel occupy peripheral zones. These roles bear different physical and cognitive risks. Fourth, the work product is itself a hazard: a partially built and cured wall whose strength evolves over hours and where workers must approach for inspection, finishing, or steel placement.

The empirical work that does exist confirms the relevance of these factors but does not resolve them. Sawicki et al. (2026) presented the first quantitative comparison of human–robot collaborative shotcrete printing against manual cast concrete in equivalent reinforced beams, observing roughly threefold higher task-level productivity and reduced physical strain (carried weight at 44%, covered distance at 37%, perceived exertion at 63% of manual reference values) . However, it also flagged the high mental demand of the operator of the robotic system as a key finding, with heart rate and blood lactate unchanged between processes. In this case, strain shifted from the body to the mind and the framework that would catch that shift does not yet exist for 3DCP.

Generic collaborative robot (i.e. cobot) frameworks have made progress on adjacent problems. For instance, PRO-MIND (Lagomarsino et al., 2024) adapts robot trajectories using heart-rate variability and frantic-movement estimates of operator stress. The largest closed-loop perceived-safety study to date, with 48 subjects on a Kinova Gen3 cobot, found that heart-rate-based speed modulation reduced productivity but did not improve perceived safety (Tusseyeva et al., 2022). In addition, continuous physiological stress models such as cStress demonstrate that lab-quality recall (88.6% on lab data) drops substantially in the field (median 72%, an 18-percentage-point gap) (Hovsepian et al., 2015). These findings transfer in principle, but the platforms, the worker roles, and the process do not.

This paper addresses these limitations by proposing a reference architecture that explicitly accounts for platform characteristics, worker roles, and process coupling in 3DCP. The design space of 3DCP HRC is organized into a three-layer taxonomy (Human, Robot, and Site) and instantiate that taxonomy as a site-wide predictive digital twin. This, in turn, acts as a continuously updated, role-aware, forward-simulating model of the print site whose outputs drive the robot's interaction with each worker present.

Related work

Standards landscape

Three regulatory frameworks bear directly on 3DCP HRC. ISO/ASTM 52939:2023 sets quality assurance requirements for additive construction but as noted, excludes EHS aspects of robotic operation from its scope (ISO/ASTM, 2023). ISO 10218-1:2025 and ISO 10218-2:2025 specify safety requirements for industrial robots and integrators respectively, with collaborative-application content (formerly ISO/TS 15066) integrated into the series. The 2025 revisions also clarify functional safety requirements and add cybersecurity provisions (ISO, 2025). The latest European Union (EU) regulation 2023/1230 on machinery, set to apply starting January 2027, introduces specific requirements for AI-based safety functions and treats software performing a safety function as a safety component subject to conformity assessment (European Parliament and Council, 2023). Together, these instruments define the regulatory envelope into which any 3DCP HRC architecture must fit, but none specifies how or which observable signals should drive that fit.

Beyond regulatory frameworks, a substantial body of research has explored the use of wearable and physiological sensing for estimating human affective and cognitive states in collaborative environments.

Psychophysiological monitoring in human–robot collaboration

The WESAD dataset (Schmidt et al., 2018) enabled standardized evaluation of multimodal stress detection from chest- and wrist-worn devices. The original benchmark reported up to 80% accuracy on a three-class problem (baseline, stress, amusement) and up to 93% on the binary stress versus non-stress task. Field studies routinely report accuracy losses relative to lab conditions. For instance, cStress, trained on a rigorous laboratory protocol with 21 participants and validated in a week-long field study with 20 participants, reported median field accuracy of 72% against self-reports compared with 90% in the laboratory (Hovsepian et al., 2015). These ceilings shape what any psychological-state-aware safety system can realistically claim.

In collaborative robotics, two threads matter for the present work. Trajectory-adaptive frameworks such as PRO-MIND (Lagomarsino et al., 2024) close a loop between operator state and robot motion. In this case, Heart Rate Variability (HRV) and frantic-movement features estimate stress and B-spline trajectories are re-optimised to balance smoothness against execution time. The empirical perceived-safety literature provides cautions. For instance, the 48-subject study by Tusseyeva et al. (2022) found that real-time speed reduction based on operator heart rate did not improve perceived safety but did reduce productivity, and that perceived safety improved with experimental habituation regardless of algorithm. The lesson is that physiological closed-loop control is not automatically safer in the user's perception and predictability matters as much as responsiveness.

Although psychophysiological monitoring enables estimation of human state in human–robot collaboration, its interpretation becomes more complex in construction environments where cognitive workload, multi-role interaction, and process constraints dominate human factors.

Human factors in construction robotics

Human factors research in construction robotics remains limited, with the field largely dominated by mechanical, materials, and process-oriented studies (Onososen and Musonda, 2023). In 3D concrete printing, Sawicki et al. (2026) show that collaborative shotcrete printing increases operator mental demand despite reduced physical strain, highlighting cognitive load as a key safety factor even when traditional physiological signals remain unchanged. Complementary simulation-based studies have explored spatial and cognitive aspects of human–robot interaction in construction, including game-based environments for analyzing interaction dynamics (Mansour et al., 2025).

Gaps and contributions

Existing cell-based collaborative robot frameworks do not transfer directly to 3D concrete printing. They assume a single operator role, a fixed and guarded workspace, localized robot hazards, and interruptible processes. On the other hand, 3D concrete printing involves multiple concurrent roles, an evolving open workspace, gantry-scale hazard envelopes, and continuous material deposition where interruptions can compromise structural integrity. Existing regulatory frameworks or data-driven HRC approaches in human–robot collaboration do not resolve these structural mismatches. Despite advances in regulatory frameworks, human state sensing, and HRC methods, there remains no unified representation that integrates human, robot, and site-level dynamics for real-time decision-making in construction environments. Physiological monitoring provides noisy estimates of cognitive and affective states that degrade in field conditions, while existing HRC approaches remain largely reactive and do not incorporate broader contextual or multi-actor site information. In parallel, construction robotics research highlights increasing cognitive workload and complex role interactions that are not captured by conventional sensing or control frameworks.

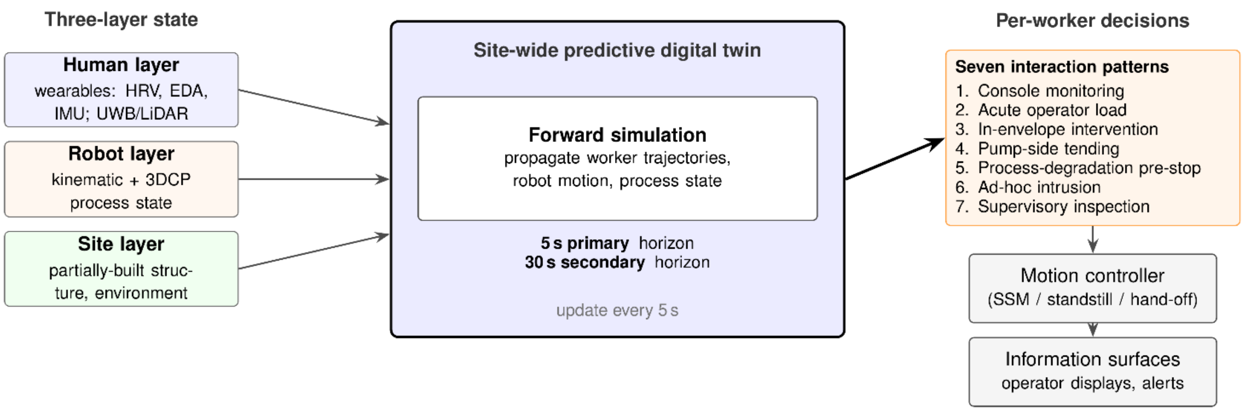

Figure 1: Proposed framework

These limitations reveal a gap between human state estimation, robot control strategies, and the operational realities of 3D concrete printing.

As such, this paper makes three contributions. First, it introduces a three-layer Human–Robot–Site taxonomy defining the required observables and state variables. Second, it develops a predictive digital-twin architecture that instantiates this taxonomy as a site-wide, continuously updated model with short-horizon forward simulation. This explicitly addresses gaps in existing safety provisions not covered by ISO/ASTM 52939 and extends beyond the assumptions of ISO 10218 collaborative robot standards. Third, it proposes a set of seven role-aware interaction patterns that map system states into actionable, standards-aligned robot behaviors for safe and auditable collaboration in 3D concrete printing. Figure 1 illustrates the proposed framework and highlights its three core components, which correspond to the contributions outlined above.

Component 1: The three-layer taxonomy

The proposed taxonomy organizes HRC-relevant observables and state variables into three layers reflecting what must be twinned to support safe collaboration. The three layers correspond to the three independently instrumented domains of the 3DCP print site and include the human counterparts, the robotic system, and the site they share, each with distinct sensing modalities, update cadences, and prediction horizons (Figure 2).

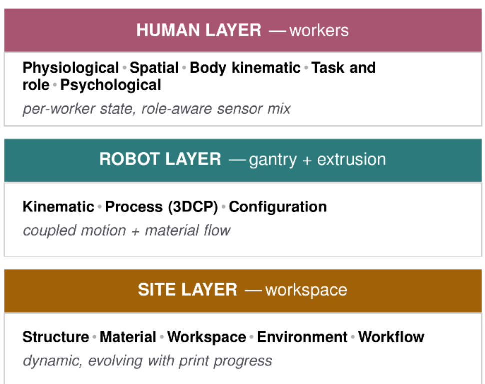

Figure 2: Proposed three-layer taxonomy

Human layer

The Human layer represents each worker present in the work envelope as an instance with five categories of state.

More specifically, the *physiological state* captures cardiovascular and electrodermal indicators relevant to acute stress and cognitive load: heart rate, heart-rate variability (RMSSD, SDNN, low-/high-frequency ratios), electrodermal activity, respiratory rate, and skin temperature. Where consented and operationally feasible, electroencephalography contributes a complementary cognitive-load and fatigue channel. In this case, industrial-grade in-helmet EEG products have demonstrated practicality at scale (e.g., over 5,000 deployed users in mining and heavy-vehicle operation), making it a credible addition to the sensor suite for operator stations rather than for all roles.

The spatial state captures three-dimensional position, velocity vector, body orientation, and a short-horizon predicted trajectory of the worker's torso and effective-arm-reach envelope. This is what enables proximity-based safety to function.

The body kinematic state, drawn from a small set of inertial measurement units (IMUs), captures three-axis acceleration, angular velocity, posture classification (standing, bent, kneeling, prone), activity classification, sudden-event detection (falls, slips, startles), and cumulative physical exertion. This layer differentiates a stationary supervisor from a rebar worker working between layers. Without it, role inference is incomplete.

On the other hand, *the task and role context* is a discrete attribute attached to each worker namely: role (operator, rebar worker, material handler, passer-by, and supervisor), current task, and experience level. More specifically, role determines which of the other state categories matter most for a given worker. For instance, a passer-by needs accurate spatial state but no physiological monitoring whereas an operator needs cognitive-load monitoring but their physical exertion is low. It is worth noting as well that the five-role list reflects the differentiation observed across current 3DCP deployments (Sawicki et al., 2026) and the framework accepts additional roles via the same role-context tag without structural change.

The inferred psychological state is a derived layer producing scalar or categorical estimates of stress, cognitive load, fatigue, and acute fear, computed from the physiological and kinematic channels above. The inference quality is constrained by the performance limits observed in field deployments, as reported in the previous section.

Robot layer

The Robot layer represents the printer as a coupled kinematic-process entity and includes several states

The kinematic and dynamic state includes: (1) position, velocity, and acceleration along the principal axes, (2) predicted trajectory along the planned toolpath, (3) current mass distribution (which on a gantry varies during extrusion as material moves through the system), (4) stopping distance under emergency braking, (5) vibration profile, and (6) instantaneous power and reaction forces. These are needed to compute speed and separation monitoring (SSM) envelopes that respect the platform's actual dynamics rather than an idealized cobot.

The process state is what makes 3DCP distinct. It includes the pump's commanded versus measured flow rate, pump-to-nozzle hose delay and rheological transit time, rheological state at the nozzle (yield stress, viscosity, thixotropic build-up), current toolpath segment, bead, and layer index, and the running gap between commanded velocity and feed rate. A robot whose kinematic state is nominal but whose process state is degrading is approaching a quality failure that may force an unplanned stop, and the safety system needs to know that.

Configuration captures setup-time invariants: bead height (H) and width (W), mix identification, and configured speed limits. These define the operating envelope within which kinematic and process state are interpreted.

Site layer

The Site layer represents the shared workspace as a dynamic structure rather than a static enclosure and includes several components.

The partially-built structure, unique to additive construction, is the most consequential novelty. It includes per-bead geometry, time since deposition (distributed across the structure as different elements were laid at different times), distributed curing state, and the nozzle's current position within the structure. A wall section deposited eight minutes ago is structurally weaker than one deposited eighty minutes ago and a worker leaning on the former is in a different hazard regime than one leaning on the latter. *On the other hand, the material delivery* tracks the pump-to-nozzle hose path and length, batch-plant state, mix supply, and material temperature. These determine when the process state at the nozzle will diverge from commanded values.

The workspace geometry tracks gantry-frame boundaries, hazard zones, sensor coverage maps (where can the proximity layer see, and where are blind spots), and the positioning of all workers and equipment within the envelope.

The *environmental state* includes ambient temperature, humidity, wind, and lighting, all of which influence both worker physiological baselines and concrete curing rates.

The *workflow and scheduling* captures the print-phase index, the schedule against which it is running, and the assignment of personnel to tasks over time. This is what enables the safety system to anticipate, rather than merely react to, role-aware interactions.

In summary, the three human, robot, and site layers are domains of content and are orthogonal to the functional decomposition of the safety controller (typically a fast proximity-tracking pipeline operating in tens of milliseconds and a slower physiological-state pipeline operating in seconds), which is well-established in the literature (Lagomarsino et al., 2024). The novelty here is not the functional split but the content schema that both pipelines need to reference.

Component 2: Predictive digital twin instantiation

A taxonomy describes but an architecture must compute. In this case, the three-layer taxonomy is instantiated as a site-wide predictive digital twin (Figure 1). The three taxonomy layers feed a continuously updated twin state, which forward-simulates worker trajectories, robot motion, and process state at 5 s primary and 30 s secondary horizons, emitting per-worker interaction-pattern decisions every 5 s. Additional details are provided in the following subsections.

Twinning fidelity across layers

Each layer requires a different update cadence and prediction horizon. Human-layer spatial and kinematic states must update at the perception subsystem's native rate (typically 10–30 Hz from UWB or LiDAR) to support proximity tracking. Human-layer physiological state updates at a slower pace consistent with the underlying signal physiology and feature-extraction windows. In this case, five seconds is a defensible primary cadence, matching the shortest meaningful HRV-feature horizon and giving the inference subsystem a stable estimate to publish. On the other hand, the robot-layer kinematic state updates at the controller's native rate (kHz at the joint level, but tens of Hz at the supervisory level relevant for HRC). The robot-layer process state updates at the pump-to-nozzle latency timescale (seconds). The site-layer partially-built structure state updates per layer or per bead, with curing models running on a slow background timer. The environmental state updates on tens-of-seconds intervals.

Developing the forward simulation mechanism

A reactive safety system catches a hazard when it has happened. A predictive system catches it before it happens which is of utmost importance in this study. Accordingly, the predictive digital twin runs a short-horizon forward simulation that propagates worker trajectories, robot motion, and process state forward under their current commands. There is a primary horizon of five seconds and another secondary horizon of thirty seconds. The horizon is bounded by the prediction error of the spatial models. Beyond five seconds, intent becomes the dominant uncertainty rather than dynamics. The simulation produces, per worker, a probability surface over near-future positions and a corresponding intersection probability with the gantry envelope and process-active zone. Figure 3 illustrates the forward simulation across the primary (5 s) and secondary (30 s) horizons. The supervisory cycle issues a new pattern decision every 5 s based on predicted intersection between worker trajectories and the gantry envelope.

Figure 3: Forward simulation across the primary (5 s) and secondary (30 s) horizons.

Updating cadence and decision cycle

The supervisory decision cycle runs at five-second cadence. At each cycle, the digital twin (1) ingests the latest layer states, (2) computes per-worker forward simulations to the primary and secondary horizons; (3) computes interaction-pattern decisions for each worker, and (4) emits supervisory references to the underlying motion controller. Faster reflexes such as emergency stops and immediate proximity violations remain on a short-loop pipeline and are not gated on the supervisory cycle.

Critically, the digital twin does not collapse its outputs into a single site-wide safety scalar. Roles have different risk profiles and different optimal interactions. A single-scalar output erases that structure. Instead, the digital twin outputs a discrete interaction pattern per worker, drawn from the seven-pattern set described in the following section.

Component 3: Role-aware interaction patterns

Seven interaction patterns were identified, each defined by who the worker is, what state the digital twin estimates, and what response the robot adopts. The seven patterns emerge from cross-tabulating the worker roles observed in current 3DCP deployments (i.e. operator, rebar worker, material handler, supervisor, passer-by) with the dominant trigger types (spatial, physiological, process-state); together they cover the operationally common cases for gantry-based 3DCP, and a deployment may add or specialize patterns without breaking the schema. Table 1 summarizes the seven patterns. Each pattern combines a role qualifier, a state qualifier from one or more taxonomy layers, and a response specification that can be mapped to an ISO 10218-1/2:2025 collaborative-application mode. The advantage of expressing safety responses as a finite pattern set, rather than as a continuous function of state, is auditability. In this case,

Table 1: Seven role-aware interaction patterns and associated information

Pattern	Role	Triggering state	Digital Twin response	ISO 10218 mode
1. Console monitoring	Operator	Cognitive-load and fatigue estimates within nominal range; large spatial separation	Adjust information presentation (alert pacing, notification suppression); no motion change	Automatic operation
2. Acute operator load	Operator	Spike in cognitive-load or fear estimate from console-side wearables	Precautionary slowdown of next path segment; pause before high-consequence action	Speed and separation monitoring (SSM)
3. In-envelope intervention	Rebar worker	Spatial intrusion within envelope during scheduled inter-layer pause	Hand-off mode: motion constrained; re-engagement gated on task-completion signal or timeout	Hand guiding / SSM
4. Pump-side tending	Material handler	Moderate proximity; process-state divergence (commanded vs measured flow)	Monitor degradation; pre-position transition to Pattern 5 if divergence persists	SSM
5. Process-degradation pre-stop	All present	Imminent quality-driven stop predicted from process state	Extend path to graceful transition (end of bead or layer); notify Patterns 3 and 4 of pending pause	Monitored standstill (planned)
6. Ad-hoc intrusion	Passer-by	Spatial intrusion into envelope; no role data available	Standard SSM response; no role-specific or physiological gating	SSM
7. Supervisory inspection	Supervisor	Permitted in-envelope presence with role qualifier; proximity to recently-deposited elements	Moderate motion; consult Site layer to flag low-strength elements on supervisor's path	SSM with site-aware caution

each pattern can be traced to the standards-defined collaborative-application types (ISO, 2025) and to a specific risk hypothesis.

Scenario-based evaluation and discussion

Given the modular structure of the proposed framework (Figure 1), the evaluation focuses on the Human-layer physiological-state inference pipeline. This includes the feasibility of end-to-end signal processing, stress classification, and calibration mechanisms using a proxy multimodal dataset. This choice is motivated by the availability of suitable datasets and the ability to conduct controlled, reproducible testing in isolation. The Robot and Site layers, as well as their integration within the predictive digital twin and interaction patterns, require system-level deployment in construction environments and are therefore reserved for future validation.

Rebar worker during inter-layer pause scenario

Consider a rebar worker who enters the print envelope at t = 0 to install reinforcement during a scheduled inter-layer pause. At the next supervisory cycle (t = 5 s), the spatial subsystem flags intrusion within the envelope. The role-context tag identifies the worker as a rebar worker and the workflow-scheduling layer confirms an inter-layer pause is in progress. The digital twin classifies this configuration as Pattern 3 (i.e. in-envelope intervention) and transitions the gantry to a hand-off mode. Motion is constrained to the SSM-compliant envelope and the supervisory cycle permits re-engagement only upon either an explicit task-completion signal from the rebar worker or the expiration of a configured timeout. The process-state subsystem continues independently. If rheology at the nozzle degrades beyond threshold during the hand-off, Pattern 5 (i.e. process-degradation pre-stop) supersedes Pattern 3, and the gantry extends path completion to a graceful transition (end of bead) rather than executing an immediate halt. After the placer signals completion, the cycle returns to Pattern 1 (i.e. console monitoring) and motion resumes. The example illustrates two framework properties: pattern selection is per worker rather than global, and pattern transitions can be process-driven rather than only space-driven.

Preliminary feasibility check

To verify that the physiological-state pipeline of the Human layer is implementable end-to-end, a preliminary benchmark was conducted using the public WESAD dataset (Schmidt et al., 2018), which provides clean multimodal physiological signals (ECG, EDA, and skin temperature) with standardized stress labels. Although the dataset was not originally collected for construction applications, its multimodal structure provides sufficient signal richness to support a proxy-based evaluation of the proposed framework. In this study, it serves as a controlled proxy for validating preprocessing, windowing, model training, threshold selection, and evaluation procedures prior to the availability of construction-specific data.

Two model classes were evaluated under a strict subject-holdout protocol (train: S2–S11, validation: S13–S14, test: S15–S17) on the binary stress versus non-stress task. The models were (1) XGBoost (Chen and Guestrin, 2016), trained on engineered window-level features (summary statistics and temporal trends), and (2) a Temporal Fusion Transformer (Lim et al., 2021), trained on sequence windows. Signals were segmented into fixed-length windows with 50% overlap. They were standardized per subject using training-set statistics. The decision threshold was selected on the validation set by maximizing the F1 score. It was then applied unchanged to the test set. Each model was evaluated under two settings: with and without per-subject calibration profiles. These profiles consist of short resting segments appended as static covariates. They correspond to the calibration mechanism prescribed by the proposed framework for operator stations. Experiments used a deliberately reduced configuration, including downsampling and a capped number of windows per subject. This configuration was chosen to support rapid iteration. The reported results therefore demonstrate pipeline functionality rather than state-of-the-art performance.

Table 2 reports the Area Under the Receiver Operating Characteristic curve (AUROC), Area Under the Precision–Recall Curve (AUPRC), and F1 on the held-out test split (11,988 windows; stress prevalence ≈ 0.31). AUROC measures the model's ability to distinguish between classes across all decision thresholds. AUPRC is more informative under class imbalance, as it focuses on performance for the positive class. F1 score is the harmonic mean of precision and recall and reflects the balance between false positives and false negatives at a chosen threshold. With calibration profiles enabled, the Temporal Fusion Transformer outperforms the tree-based baseline. With profiles disabled, the sequence model performance collapses while XGBoost remains stable. This indicates stronger reliance of the sequence model on the personalization pathway prescribed by the framework. To ensure comparability, train/validation/test splits and class distributions were kept identical across configurations. The profiles-off configuration was rerun and reproduced identical metrics. Threshold selection was performed exclusively on the validation set.

Table 2: WESAD pipeline feasibility check on the held-out test split

Setting	Model	AUROC	AUPRC	F1
Profiles on	XGBoost	0.712	0.542	0.498
Profiles on	TFT	0.768	0.683	0.576
Profiles off	XGBoost	0.712	0.542	0.498
Profiles off	TFT	0.548	0.354	0.302

These results demonstrate that the Human-layer inference pipeline can support stable stress classification and calibration under controlled conditions, representing an initial step toward operationalizing the proposed framework.

The seven-pattern model, role-aware response logic, and process-coupled patterns (4 and 5) remain to be evaluated in deployment settings.

Conclusion, limitations, and future work

This paper has presented a three-layer taxonomy (Human, Robot, Site), an instantiating site-wide predictive digital twin for human–robot collaboration in 3D concrete printing, seven role-aware interaction patterns, and a scenario-based evaluation of the Human-layer physiological-state inference pipeline. The proposed taxonomy structures the design space of observables and state variables, while the digital twin operationalizes this structure into a continuously updated model. This model supports seven role-aware interaction patterns mapped onto ISO 10218-1/2:2025 collaborative-application modes. The framework is platform-agnostic across gantry, robot-arm, and crane-based 3DCP systems and is intended as a reference architecture for future deployment and standardization. The work is motivated by a regulatory gap, as ISO/ASTM 52939:2023 excludes the EHS aspects of robotic operation, and by an empirical gap in existing collaborative-robot frameworks, which do not account for process-coupled hazards or multi-role construction environments.

While the paper presents an innovative framework, it does not constitute a deployed system. The main limitation is that the proposed interaction patterns and layered digital twin have not yet been validated in integrated real-world 3D concrete printing deployments. The Human-layer evaluation demonstrates feasibility under controlled conditions, but full system-level validation across Human, Robot, and Site layers remains to be conducted.

Future work will proceed in three directions. First, system-level deployment and validation will be carried out on an operational gantry-based 3DCP platform, including quantitative assessment of false positives and false negatives under scripted hazard scenarios. Second, role inference mechanisms will be developed to enable automatic classification of worker roles from kinematic and contextual signals, reducing reliance on manual role assignment. Third, the interaction pattern set will be formalized as a normative contribution aligned with the ISO 10218 series, alongside engagement with ISO/ASTM TC261/F42 to support future revisions of ISO/ASTM 52939.

Acknowledgements

The presented work was supported by the American University of Beirut (AUB) Vertical Integration Projects (VIP) Program. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the VIP Program or its stakeholders.

References

Bartneck, C., Kulić, D., Croft, E. and Zoghbi, S. (2009) Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. International Journal of Social Robotics, 1(1), pp.71–81.

Buswell, R.A., Leal de Silva, W.R., Jones, S.Z. and Dirrenberger, J. (2018) 3D printing using concrete extrusion: A roadmap for research. Cement and Concrete Research, 112, pp.37–49.

Chen, T. and Guestrin, C. (2016) XGBoost: A scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. New York, NY, USA, ACM, pp.785–794.

COBOD International (2024) First 3D Printed Social Housing Project Built to New ISO/ASTM Standard Completed 35% Faster Than With Conventional Methods. Press release, 17 December 2024.

European Parliament and Council (2023) Regulation (EU) 2023/1230 of 14 June 2023 on machinery and repealing Directive 2006/42/EC. Official Journal of the European Union, L 165, 29 June 2023, pp.1–102.

Hovsepian, K., al'Absi, M., Ertin, E., Kamarck, T., Nakajima, M. and Kumar, S. (2015) cStress: Towards a gold standard for continuous stress assessment in the mobile environment. In: Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15). New York, NY, USA, ACM, pp.493–504.

ISO (2025) ISO 10218-1:2025 — Robotics — Safety requirements — Part 1: Industrial robots and ISO 10218-2:2025 — Robotics — Safety requirements — Part 2: Industrial robot applications and robot cells. Geneva, International Organization for Standardization.

ISO/ASTM (2023) ISO/ASTM 52939:2023 — Additive manufacturing for construction — Qualification principles — Structural and infrastructure elements. Geneva, International Organization for Standardization.

Lagomarsino, M., Lorenzini, M., De Momi, E. and Ajoudani, A. (2024) PRO-MIND: Proximity and Reactivity Optimisation of robot Motion to tune safety limits, human stress, and productivity in INDustrial settings. arXiv preprint arXiv:2409.06864. (Forthcoming, IEEE Robotics and Automation Letters.)

Lim, B., Arık, S.Ö., Loeff, N. and Pfister, T. (2021) Temporal fusion transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting, 37(4), pp.1748–1764.

Mansour, A., Abou Khalil, A. & Khoury, H.M. (2025) Toward spatial and cognitive human–robot interaction: A game-based simulation for concrete 3D printing in construction. In: Proceedings of the Creative Construction Conference (CCC). Zadar, Croatia, IAARC, pp. 750–755.

Marvel, J.A. (2013) Performance metrics of speed and separation monitoring in shared workspaces. IEEE Transactions on Automation Science and Engineering, 10(2), pp.405–414.

Onososen, A.O. and Musonda, I. (2023) Research focus, trends and outlook on construction robotics for human-robot teams: A scientometric review. Journal of Engineering, Design and Technology, 21(2), pp.502–526.

Panda, B., Paul, S.C., Mohamed, N.A.N., Tay, Y.W.D. and Tan, M.J. (2018) Measurement of tensile bond strength of 3D printed geopolymer mortar. Measurement, 113, pp.108–116.

Roussel, N. (2018) Rheological requirements for printable concretes. Cement and Concrete Research, 112, pp.76–85.

Sawicki, B., Düking, P., Placzek, G., Masur, L., Dörrie, R., Schwerdtner, P. and Kloft, H. (2026) Human–robot collaboration in digital fabrication with concrete: quantifying productivity and psychophysiological strain of human workers. Construction Robotics, online first.

Schmidt, P., Reiss, A., Duerichen, R., Marberger, C. and Van Laerhoven, K. (2018) Introducing WESAD, a multimodal dataset for wearable stress and affect detection. In: Proceedings of the 20th ACM International Conference on Multimodal Interaction (ICMI). Boulder, CO, USA, ACM, pp. 400–408.