AI-Based Structural Performance Prediction for Robotic Concrete Extrusion in Complex Building Geometries

Willem Shupe
Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
willemshupe92@ku.edu

Tane Costello
School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
tanecostello00@uc.edu

Engla Sjolander
School of Computing, Clemson University, Clemson, SC, USA.
engla732@clemson.edu

Serafin Lemus
Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
serafinl@buffalo.edu

Abstract

Robotic concrete extrusion can realize non-prismatic walls, branching columns, and cantilevered envelopes, yet the structural outcome of a print is decided before the first layer is placed: it follows from the interaction of fresh-state rheology, time-dependent structuration, interlayer delay imposed by the toolpath, and the eccentricity that geometric complexity introduces into the load path. Existing buildability criteria reduce this interaction to a scalar for a monotonically rising vertical wall, and therefore misstate the margin in geometries where support is non-local and where overhang shifts the centroid of the deposited stack. This paper specifies a geometry-aware sequential surrogate framework that predicts layer-resolved structural margins directly from a fabrication plan. The contributions are fourfold: the prediction target is reformulated as a time-indexed stability margin rather than a single failure height; the print is represented as a heterogeneous graph whose typed edges encode toolpath order, vertical support, and lateral continuity, so that non-local support relations in complex geometry remain expressible; a physics-informed multi-task objective anchors the learned mapping to Bingham extrusion, eccentricity-induced bending, and elastic layer buckling; and distribution-free conformal calibration supplies the uncertainty term needed for an explicit admissibility rule during print-plan screening. The manuscript develops the governing relations, the network architecture and loss composition, the identifiability conditions under which fresh-state parameters can be inferred, and a domain-specific evaluation protocol that distinguishes corner overhang, thin-wall buckling, and cold-joint error modes.

Keywords

robotic concrete extrusion; buildability prediction; physics-informed machine learning; graph neural networks; geometric complexity; conformal uncertainty

1. Introduction

Extrusion-based robotic fabrication of concrete has progressed from isolated demonstrations to a research programme with identifiable material, process, and control requirements. Roadmaps for the technology locate the central difficulty not in the deposition mechanism itself but in the coupling between machine motion, material rheology, and the geometry being produced [1]. Surveys of additive manufacturing in construction reach a similar conclusion, noting that deposition rate, nozzle geometry, and path planning act jointly on the quality of the deposited filament and therefore cannot be optimized in isolation [2].

Rheological requirements for printable concretes establish the first constraint: a material must be pumpable through a delivery line, extrudable through a nozzle, and simultaneously capable of retaining shape once the driving pressure is removed [3]. Shape retention is a transient property. The structural built-up problem, formulated analytically for a vertical wall printed at a constant vertical rate, shows that the layer stack must carry an increasing self-weight with a strength that increases only through rest-time structuration, so that failure occurs at a height determined by the ratio of the two rates [4]. Numerical treatment of early-age mechanical behaviour confirms that this competition, rather than the absolute strength of the hardened material, governs the failure height [5]. A consolidated review of the underlying physics identifies flocculation, desorption, hydration, and the drainage of the lubricating layer at the wall slip boundary as the mechanisms that set the time scales over which the competition is resolved [6].

Analytical buildability models derived from this picture are conservative and inexpensive, but they assume a monotonically rising vertical wall with a uniform interlayer interval [7]. Complex building geometries violate both assumptions. Cantilevered and overhanging elements have been analysed through a masonry analogy in which the stability of the printed stack is governed by the sequence of eccentricities accumulated across successive layers rather than by vertical stress alone [8]. The distinction matters because fresh concrete tolerates compressive stress far better than the tensile stress that eccentricity generates on the inner face of an overhang. Hardened properties compound the problem: printed concrete is direction-dependent, with interlayer and interstrip interfaces acting as weak planes [9], and the anisotropy is measurable in both flexural and compressive response [10]. Robotic slip-form casting of façade elements shows that a moving formwork can extend the admissible geometric envelope by carrying part of the fresh-state load, which indicates that the load path during fabrication is itself a design variable [11]. Reviews of printing trends in construction observe that the geometric freedom claimed for the process is bounded in practice by buildability and by the absence of prediction tools that operate on the fabrication plan rather than on a material specimen [12].

The gap addressed here is therefore not the absence of machine learning in construction, nor the absence of fresh-state mechanics, but the absence of a formulation that maps a complete fabrication plan, including its geometric complexity, onto layer-resolved structural margins with quantified uncertainty. Purely data-driven regressors trained on buildability outcomes extrapolate poorly because the input space of geometries is unbounded and the failure surface is non-smooth at geometric transitions. High-fidelity finite-element simulation of the fresh state is more faithful but costs minutes to hours per plan and requires constitutive parameters that are rarely available on site, which precludes its use inside a plan-screening or feed-rate-optimization loop.

This paper specifies a framework that occupies the intermediate position. The prediction target is reformulated as a time-indexed margin at every toolpath segment, so that the output is a spatial field of admissibility rather than a single number. The print is represented as a heterogeneous graph whose typed edges encode toolpath order, vertical support, and lateral continuity; this representation preserves non-local support relations that a purely sequential model cannot express. A physics-informed multi-task objective anchors the learned mapping to the Bingham extrusion relation, to eccentricity-induced bending, and to elastic layer buckling. Conformal calibration then converts the residual distribution into a per-segment interval, which supports an explicit admissibility rule for plan screening. The remainder of the paper develops the governing relations and the formal prediction target, the framework architecture, the identifiability and calibration conditions, and the evaluation logic together with its domain-specific error modes.

2. Governing Mechanisms and Problem Formulation

2.1 Extrusion rheology and deposition state

Fresh printable concrete is conventionally described as a yield-stress fluid, and over the shear-rate range imposed in a delivery line and nozzle its response is well approximated by a Bingham constitutive relation,

$$\tau = \tau_0 + \mu\,\dot{\gamma}, \dot{\gamma} > 0$$

where τ is the shear stress, τ_0 the static yield stress, μ the plastic viscosity, and $\dot{\gamma}$ the shear rate. The relation is piecewise in the sense that no deformation occurs below τ_0, which is why the material can retain a deposited shape. Two consequences follow for prediction. First, the wall shear rate in the nozzle scales approximately with the mean velocity divided by the nozzle diameter, so that the extrusion rate and the print speed are not independent controls; the pair determines the shear history to which the material is subjected immediately before deposition. Second, τ_0 is not a constant of the mixture but a function of the rest time since the material last experienced shear above the yield threshold. This thixotropic recovery is the mechanism that allows a lower layer to support the weight of layers placed above it.

2.2 Time-dependent strength and layer loading

The evolution of the structural contribution of the material is captured by a structuration law in which the strength available at the base of a layer grows toward an asymptote,

$$S(t) = S\infty\big(1 - \exp(-kt)\big)$$

where S(t) is the available resistance at rest time t measured from deposition, $S\infty$ is the asymptotic structural resistance of the mixture, and k is a structuration rate with units of inverse time. For rest times short compared with 1/k, the exponential can be expanded and the law reduces to the linear thixotropic form $S(t) \approx S\infty\ k\ t$, which recovers the constant-rate build-up assumed in analytical buildability models [4], [7]. The exponential form is retained here because interlayer intervals in complex geometries are frequently long at geometric discontinuities, where the nozzle must decelerate, retract, or travel to a distant region, and in that regime the linear approximation overestimates the strength gained.

Loading is accumulated by the layers above. For a stack discretized into N layers of deposited height h_j and fresh density ρ, the vertical stress acting at the base of layer i at the instant the final layer is placed is

$$\sigma_{v,i} = \rho\, g\, \Sigma_{j=i+1}{}^{N} h_j$$

with g the gravitational acceleration. The index j runs over all layers deposited above layer i. In a vertical wall this sum is simply the height of material above the layer, and the critical layer is always the lowest one. In a complex geometry the sum is not the relevant quantity, because the load path is not vertical: material may be supported by a distant column, by a branching junction, or by a substrate. This is the first reason a scalar buildability index fails, and it motivates a representation in which support is expressed as a relation rather than as a height.

2.3 Eccentricity, bending, and geometric complexity

Where the toolpath of a layer is offset horizontally from the centroid of its supporting footprint, the weight transmitted through that layer acquires an eccentricity. For an overhang produced at an angle α to the vertical with a layer height h, consecutive deposits shift laterally by approximately h tan α, and the offset accumulates along the height of the supported stack. Denoting by e_i the horizontal offset between the centroid of the material above layer i and the centroid of the footprint that supports layer i, the stress on the extreme fibre of a rectangular section of thickness h_i follows from the superposition of axial load and bending moment,

$$\sigma_i = \sigma_{v,i}(1 + 6e_i / h_i)$$

which is the standard result obtained by combining P/A with Mc/I for M = P e_i, c = h_i/2, and I = b h_i^3/12 for unit width b. The amplification factor $6e_i/h_i$ is the quantitative expression of the geometric penalty: an eccentricity of one twelfth of the layer thickness already doubles the extreme-fibre stress. Because the same expression governs the tensile side of the section, and because fresh concrete has a tensile resistance that is small compared with its compressive resistance, the margin on the tensile side frequently controls admissibility in overhanging regions even when the compressive margin appears comfortable.

A second geometric failure mode is elastic buckling of a slender printed element. A thin wall or a free-standing printed rib can lose stability before the material reaches its compressive capacity, and the critical stress for a plate strip of thickness h_i and in-plane width b_i takes the form

$$\sigma cr,_i = C_b\, E(t_i)(h_i/b_i)^2$$

where C_b is a coefficient that absorbs edge restraint, Poisson effects, and the ratio of the buckle half-wavelength to the wall width, and $E(t_i)$ is the secant modulus of the material at the age of layer i. Because $E(t_i)$ grows with rest time in the same manner as $S(t_i)$, the buckling margin of a lower layer improves as the print proceeds, whereas the applied stress grows; the two effects are opposed, and the minimum of the margin over the height need not occur at the base. Classical stability theory supplies the structural forms on which this expression rests [14], while the age dependence of the elastic modulus is the standard maturity-type relation used for early-age concrete [13]. The coefficient C_b cannot be derived in closed form for the geometries of interest, and its effective value depends on restraint from adjacent material, which is one quantity the learned model is expected to absorb from data rather than to assume.

2.4 The prediction target

Combining the three mechanisms gives the quantity the framework is built to predict. Define the layer margin ratio as the applied extreme-fibre stress normalized by the smaller of the two available resistances,

$$M_i = \sigma_i / min(S(t_i)/F_s, \sigma cr,_i)$$

where F_s is a design factor applied to the structuration resistance. Admissibility of a segment requires $M_i < 1$ at every layer above it and at every instant of the print. The formulation is deliberately local in space and global in time: the margin at a given segment depends on everything deposited above it and on the schedule that governed when each of those deposits occurred. It is this dependence on schedule and support topology, rather than on material alone, that a surrogate must learn.

3. Geometry-Aware Sequential Surrogate Framework

3.1 Print representation as a heterogeneous graph

The fabrication plan, comprising the toolpath, the extrusion rate schedule, and the layer decomposition, is discretized into toolpath segments of approximately constant arc length. Each segment becomes a node carrying a feature vector that includes its position and orientation, the local toolpath curvature, the extrusion rate, the print speed, the layer index, the deposition timestamp, the material age at deposition, and the time elapsed since the supporting material beneath it was deposited. Curvature is included because it controls filament overlap and the effective deposited width at corners, which in turn sets the local thickness h_i that appears in both the bending and buckling relations.

Edges are typed. Sequential edges connect consecutive segments along the toolpath and carry the process history: pump dynamics, acceleration, and the accumulated shear history. Vertical edges connect each segment to the segments of the layer beneath that geometrically support it, identified by projecting the segment footprint onto the plane of the layer below and retaining the overlap above a threshold. Lateral edges connect adjacent segments within a layer and encode the continuity of the printed strip. The vertical edges are the critical design choice. In a vertical wall they reduce to a one-to-one chain and the graph degenerates to a sequence, but in a branching column, a cantilever, or a bridge span the support relation is many-to-many and non-local in print order, and a model that only sees the deposition sequence cannot represent it. This is the structural reason for preferring a graph representation over a purely sequential one.

3.2 Fresh-state state propagation

Rather than treating rheological parameters as static inputs, the framework propagates a fresh-state vector along the print. At deposition, a node inherits the mixture state produced by the shear history of the delivery line and nozzle. After deposition, the state evolves by structuration according to the law of Section 2.2, so that the yield stress and the structural resistance at any subsequent instant are deterministic functions of the elapsed rest time and of the mixture parameters. The propagation supplies, for every node and every time step of interest, the quantities $S(t_i)$ and $E(t_i)$ that appear in the margin ratio.

The propagation step is where the framework interfaces with in-line instrumentation. Pressure at the pump and at the nozzle inlet, together with motor torque on the extruder, provide indirect observations of the mixture state that can be assimilated into the structuration parameters as the print proceeds. The assimilation is a low-dimensional filtering problem, and its output is a posterior over the structuration rate and the reference yield stress rather than a point estimate. Carrying that posterior forward is what allows the final prediction interval to reflect material variability and not only model error.

3.3 Encoder architecture

Two encoders operate in parallel and are fused before the prediction heads. A recurrent encoder processes the toolpath sequence in deposition order, supplying the process-history context that determines the shear and rest state at each node. Gated recurrent units or long short-term memory cells are appropriate because the relevant dependence is over variable-length intervals with irregular spacing [18]. A graph encoder performs message passing over the typed graph, updating each node from its neighbours,

$$h_v^{(l+1)} = \varphi\left(h_v^{(l)}, \Sigma_u \psi\left(h_u^{(l)}, e_{uv}\right)\right)$$

where $h_v^{(l)}$ is the hidden state of node v at layer l, e_{uv} is the edge feature of the typed edge from neighbour u, ψ is a message function, φ an update function, and the sum runs over the neighbourhood of v, with separate parameter sets for sequential, vertical, and lateral edge types [17]. The two encoders are complementary by construction: the recurrent path answers when the material arrived and in what state, the graph path answers what is above it and how the load reaches the support.

Fusion is by concatenation followed by a small gated network that learns to weight the two representations per node. The gate is conditioned on the local geometric descriptors, so that regions with strong non-local support rely more heavily on the graph branch while regions that are essentially vertical rely on the sequential branch. This conditioning is preferable to a fixed weighting because the relative informativeness of the two views is itself a function of geometry.

3.4 Multi-task heads and physics-informed objective

Three heads share the fused representation. The first predicts the layer margin ratio M_i at each node, which is the primary output. The second predicts the residual vertical deformation of each layer, which supports serviceability assessment and provides an auxiliary signal that is more densely observable in physical build tests than failure itself. The third predicts the depth to failure for the plan as a whole, providing a global scalar that is directly comparable with analytical buildability criteria.

Supervision is sparse; physics constraints supply the missing structure. The total objective combines a data term, a physics residual term, and a parameter regularizer,

$$L = L_d + \lambda_p\, L_p + \lambda_r \|\theta\|^2$$

where L_d is the supervised loss on whichever labels are available at a node or plan, L_p aggregates the residuals of the constitutive, equilibrium, and stability relations evaluated on the predicted states, θ denotes the network parameters, and λ_p and λ_r are weighting coefficients. The physics residual includes the requirement that the applied stress recovered from the predicted margin be consistent with the vertical load computed from the predicted deposition heights, that the margin decrease monotonically as overlying material is added at fixed rest time, and that the predicted critical height for a vertical wall agree with the analytical bound in the limit where the geometry reduces to a wall. Embedding such relations as soft penalties rather than as hard architectural constraints preserves the ability of the model to represent the effective restraint coefficient C_b, which is not known in closed form, while preventing solutions that violate the mechanisms the data cannot constrain. The broader rationale for combining scientific knowledge with learned components in engineering systems applies directly here, since the constraints encode conservation and stability requirements that hold irrespective of the training distribution [16]. Physics-informed training of this general

form has been demonstrated to improve generalization in regimes where data are scarce and the governing equations are known [15].

3.5 Uncertainty quantification and the admissibility rule

A point prediction of the margin is insufficient for a decision that governs whether a plan is printed. Two uncertainty sources are distinguished: parameter uncertainty from the fresh-state posterior, and model discrepancy from the surrogate approximation. The first is propagated by sampling the structuration posterior and re-evaluating the deterministic stress relations, which is inexpensive because the relations are closed-form. The second is estimated by a deep ensemble or by a Gaussian process head on the fused representation, both of which are standard and well characterized in the small-data regime typical of construction research [19].

The ensemble or process variance is then calibrated by split conformal prediction. Given a calibration set of M plans with observed or simulated margins y_n and predictions $\hat{y}_n$, the calibration quantile is

$$\hat{q} = Quantile_{1-\delta}(\{|y_n - \hat{y}_n|\}_{n=1}^{M})$$

where δ is the target miscoverage rate and $\text{Quantile}_{1-\delta}$ denotes the empirical $(1-\delta)$ quantile of the absolute residuals [20]. The resulting interval $\hat{y} \pm \hat{q}$ carries a marginal coverage guarantee under exchangeability of calibration and test plans. The admissibility rule is then: a segment is admissible only if $\hat{y}_i + \hat{q} \leq \min(S(t_i)/F_s, \sigma cr,_i)$ for all layers above it. The rule is deliberately conservative on the tensile side, since the margin ratio already amplifies eccentricity, and asymmetric intervals are used where the residual distribution is skewed.

The exchangeability assumption deserves emphasis because it is the point at which the guarantee is most likely to be violated in practice. Plans drawn from a new geometric family, or printed with a new mixture batch, are not exchangeable with a calibration set dominated by a different family. The defensible response is stratification: calibration sets are partitioned by geometric class and by mixture family, and a plan is calibrated against the stratum it most closely resembles. Where no stratum is populated, the conformal interval is replaced by the wider ensemble interval and flagged for physical build verification.

4. Calibration, Identifiability, and Data Regimes

The framework contains parameters that cannot be observed directly during a print. The structuration rate k, the asymptotic resistance S_∞, and the age-dependent modulus E(t) are inferred from a small number of offline measurements, typically a squeegee or slow-penetration test for the yield stress evolution and a early-age compression series for the modulus, and then updated by assimilation during the print. The degree-of-hydration framing provides the link between the maturity of the mixture and its mechanical properties and is the natural basis for the age-dependent modulus used in the buckling relation [13].

Identifiability is the principal methodological risk. Two distinct parameter combinations can produce the same observed failure height in a vertical wall, since an increase in structuration rate can be offset by a decrease in the asymptotic resistance. The degeneracy is broken by the geometry: an overhanging region loads the material eccentrically and therefore probes the tensile side of the response, while a thin free-standing rib probes the buckling relation, and the two probe different combinations of the parameters. A training set that spans these geometric classes is consequently more informative per plan than a larger set of vertical walls, and the design of the label set should reflect this rather than maximize plan count.

Labels arise from two sources. Calibrated finite-element simulations of the fresh state, using a time-dependent yield criterion, supply dense per-node margins at a cost that is acceptable offline. Physical build tests supply sparse labels at the plan level, typically the height and location of the first observed instability. The two sources have different bias structures: the simulation inherits the constitutive assumptions, in particular the treatment of the interlayer interface as a cohesive plane, while the physical test inherits measurement noise in the layer height and in the actual rest intervals. A model trained on the mixture must carry a source indicator so that the conformal calibration can be performed separately for each, otherwise the interval width will reflect the mixture of biases rather than either one.

Dimensional analysis reduces the data requirement further. The margin ratio of Section 2.4 depends on the ratio of gravitational loading to structural resistance, on the geometric ratios e_i/h_i and h_i/b_i, and on the dimensionless rest time $k\,t$. Presenting these groups as explicit inputs, in addition to the raw features, allows the network to share statistical strength across material batches and building scales and reduces extrapolation error when a plan is outside the geometric range of the training set but inside the range of the dimensionless groups.

5. Evaluation Logic and Domain-Specific Error Modes

Evaluation must be designed around the decision the model supports, not around aggregate regression accuracy. Four measures are appropriate. The first is the per-node margin error normalized by the local available resistance, which is scale-free and comparable across geometries. The second is the depth-to-failure error, defined as the difference between the predicted and reference layer index at which the margin first exceeds unity; this measure penalizes errors at the critical layer far more than errors elsewhere and therefore aligns with the consequence structure of the problem. The third is the classification performance of the admissibility rule, reported with asymmetric costs because releasing an inadmissible plan is more damaging than rejecting an admissible one. The fourth is empirical coverage of the conformal interval by geometric stratum, which verifies that the guarantee survives the stratification.

Credible baselines are required for any claim of improvement. The analytical buildability bound of Section 2.2, evaluated with the plan's mean interlayer interval, is the first baseline and is expected to be accurate for vertical regions and optimistic for overhangs. A coarse finite-element model of the fresh state, run at a resolution that fits within the same time budget as the surrogate, is the second. A gradient-boosted regressor on hand-crafted geometric and process features is the third, and it establishes whether the graph and sequence structure contributes beyond feature engineering [21]. A sequence-only recurrent model without the graph branch is the fourth, and it isolates the contribution of non-local support representation. Reporting against all four separates the effect of representation from the effect of model capacity.

Three error modes are specific to this domain and should be reported separately rather than averaged away. The first occurs at concave corners and junctions, where the projection used to define vertical support is ambiguous because a segment overlaps two or more supporting regions with different rest times. The model must then learn a mixture of supporting stiffnesses, and the residual at these nodes is typically the largest in the print; the appropriate treatment is to expose the support overlap fractions as explicit edge features rather than to force a single support assignment. The second occurs in thin free-standing elements, where the buckling term dominates and where the effective restraint coefficient depends on adjacent material that may or may not have been printed at the time of loading. The third

occurs at geometric discontinuities, where the nozzle must decelerate or travel, producing interlayer intervals that are long compared with the mean; the structuration law is then in its nonlinear regime and the linear build-up approximation embedded in analytical baselines is systematically wrong. Because the framework evaluates the structuration law exactly at the observed interval, it is expected to be most advantageous precisely in these regions, and the evaluation should be structured to make that advantage visible or to refute it.

6. Technical Discussion

The framework occupies a specific position in the fidelity-latency trade-off. Finite-element simulation of the fresh state resolves the stress field and the interface behaviour but costs minutes to hours per plan and requires constitutive parameters that are themselves uncertain. The surrogate reduces inference to a single forward pass over a graph with tens of thousands of nodes, which is compatible with the inner loop of a toolpath optimizer that varies print speed, extrusion rate, layer height, and decomposition order. The value of the surrogate is therefore not that it replaces simulation but that it makes plan-level search tractable, with simulation reserved for the final verification of the selected plan.

A second implication concerns process control. Because the output is a per-segment margin rather than a per-plan scalar, the model can be queried for the maximum admissible print speed at each segment while holding the margin at a target value. This converts a geometric design constraint into a spatially varying feed-rate schedule, which is a natural interface to the robot controller and complements approaches in which the moving formwork itself carries part of the fresh-state load [11]. The two strategies are complementary: formwork support relaxes the margin in regions where it is applied, and the surrogate quantifies how much relaxation is available.

Full-scale demonstrators, such as printed bridges and structural elements, provide the ultimate test of any prediction framework because they expose the interaction of buildability with hardened anisotropic behaviour at the structural scale [22], [9]. The margin formulation of this paper addresses the fresh-state problem; the hardened-state performance of a complex geometry depends additionally on interface orientation relative to the principal stress directions, which is a separate prediction problem that shares the graph representation and the toolpath-derived features but requires a different label source.

The main limitations follow from the construction of the method. The coverage guarantee of the conformal interval is marginal and conditional only under exchangeability, and geometric novelty is precisely the condition under which exchangeability fails. The physics residual constrains the solution only to the extent that the assumed relations are correct, and the effective restraint coefficient in the buckling relation remains an empirical quantity that must be learned. The graph construction cost scales with the product of segment count and support overlap, which for a full building may require hierarchical coarsening. Finally, the framework's predictions are only as good as the fresh-state parameter posterior, and the assimilation of in-line pressure data into that posterior is a modelling problem in its own right.

7. Conclusion

Complex building geometries expose a weakness in the way structural performance is predicted for robotic concrete extrusion. The prediction problem is not the determination of a single failure height for a rising wall, but the determination of a margin field over the toolpath, indexed by the schedule of deposition, in which geometric overhang converts vertical load into bending and in which support relations are non-local in print order. This paper has

reformulated the prediction target accordingly, expressed the eccentricity-induced stress amplification and the competing elastic buckling mode as explicit relations, and specified a surrogate framework in which a recurrent encoder captures process history, a typed graph encoder captures support topology, and a physics-informed multi-task objective anchors both to the governing mechanisms. Conformal calibration supplies the uncertainty term required for an admissibility rule that can screen a fabrication plan before a machine is committed to it.

The methodological contribution is the alignment between representation and mechanism: the graph structure exists because support in complex geometry is not sequential, the structuration law is evaluated exactly at the observed interlayer interval because geometric discontinuities produce long rest times, and the physics residual is soft because the effective restraint coefficient is not known in closed form. The resulting framework is transparent about its own failure conditions, which are geometric novelty, thin-element restraint, and material batch shift, and it defines the stratified calibration and separate error reporting needed to detect them. Extending the same representation from fresh-state margins to hardened anisotropic performance would provide a single predictive substrate spanning the full fabrication-to-service life of a printed element.

References

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

2. Bos, F. P., Wolfs, R. J. M., Ahmed, Z. Y., & Salet, T. A. M. (2016). Additive manufacturing of concrete in construction: Potentials and challenges of 3D concrete printing. Virtual and Physical Prototyping, 11(3), 209-225.

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

4. Perrot, A., Rangeard, D., & Pierre, A. (2016). Structural built-up of cement-based materials used for 3D-printing extrusion techniques. Materials and Structures, 49(4), 1213-1220.

5. Wolfs, R. J. M., Bos, F. P., & Salet, T. A. M. (2018). Early age mechanical behaviour of 3D printed concrete: Numerical modelling and experimental testing. Cement and Concrete Research, 106, 103-116.

6. Mechtcherine, V., Bos, F. P., Perrot, A., Leal da Silva, W. R., Nerella, V. N., Fataei, S., Wolfs, R. J. M., Sonebi, M., & Roussel, N. (2020). Extrusion-based additive manufacturing with cement-based materials: Production steps, processes, and their underlying physics: A review. Cement and Concrete Research, 132, 106037.

7. Kruger, J., Zeranka, S., & van Zijl, G. (2019). 3D concrete printing: A lower bound analytical model for buildability performance of extrusion-based processes. Automation in Construction, 106, 102904.

8. Carneau, P., Mesnil, R., Roussel, N., & Baverel, O. (2020). Additive manufacturing of cantilever: From masonry to concrete 3D printing. Automation in Construction, 116, 103184.

9. Le, T. T., Austin, S. A., Lim, S., Buswell, R. A., Law, R., Gibb, A. G. F., & Thorpe, T. (2012). Hardened properties of high-performance printing concrete. Cement and Concrete Research, 42(3), 558-566.

10. Panda, B., Paul, S. C., & Tan, M. J. (2017). Anisotropic mechanical performance of 3D printed fiber reinforced sustainable construction material. Materials Letters, 209, 146-149.

11. Yu, L., Luo, D., & Xu, W. (2018). Dynamic robotic slip-form casting and eco-friendly building façade design. In Robotic fabrication in architecture, Art and Design (pp. 421-433). Cham: Springer International Publishing.

12. Tay, Y. W. D., Panda, B., Paul, S. C., Noor Mohamed, N. A., Tan, M. J., & Leong, K. F. (2017). 3D printing trends in building and construction industry: A review. Virtual and Physical Prototyping, 12(3), 261-276.

13. De Schutter, G., & Taerwe, L. (1996). Degree of hydration-based description of mechanical properties of early age concrete. Materials and Structures, 29(6), 335-344.

14. Bazant, Z. P., & Cedolin, L. (1991). Stability of structures: Elastic, inelastic, fracture and damage theories. Oxford University Press.

15. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 378, 686-707.

16. Willard, J., Jia, X., Xu, S., Steinbach, M., & Kumar, V. (2022). Integrating scientific knowledge with machine learning for engineering and environmental systems. ACM Computing Surveys, 55(4), 1-37.

17. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In Proceedings of the International Conference on Learning Representations.

18. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.

19. Rasmussen, C. E., & Williams, C. K. I. (2006). Gaussian processes for machine learning. MIT Press.

20. Angelopoulos, A. N., & Bates, S. (2021). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. arXiv preprint arXiv:2107.07511.

21. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5-32.

22. Salet, T. A. M., Ahmed, Z. Y., Bos, F. P., & Wolfs, R. J. M. (2018). Design of a 3D printed concrete bridge by testing. Virtual and Physical Prototyping, 13(3), 222-236.