AI-Driven Print-Path Generation for Scaffold-Free Robotic 3D Concrete Printing of Vaulted Structures

Abdallah Kamhawi, Dr. Mania Aghaei Meibodi, Dr. Jesse Capecelatro

DART Laboratory, Taubman College of Architecture and Urban Planning
Capecelatro Research Group, Department of Aerospace Engineering, Department of Mechanical Engineering

DART
Digital Architecture Research & Technologies

Capecelatro Research Group

01 Problem Statement

Scaffold-free robotic 3D concrete printing of vaulted structures requires every intermediate construction state to remain self-supporting under its own weight. The fabrication system used in this research (Figure 1) is a two-component (2K) robotic extrusion setup comprising a KUKA KR120 R2700 HA 6-axis industrial robot with an inline two-component mixing head, a MAI MULTIMIX-3D concrete pump, and an accelerator dosing unit. The 2K system enables programmable early-age kinetics: accelerator dosage, flow rate, and tool center point (TCP) speed are independently actuated and jointly govern structural build-up, bead geometry, and interlayer bonding.

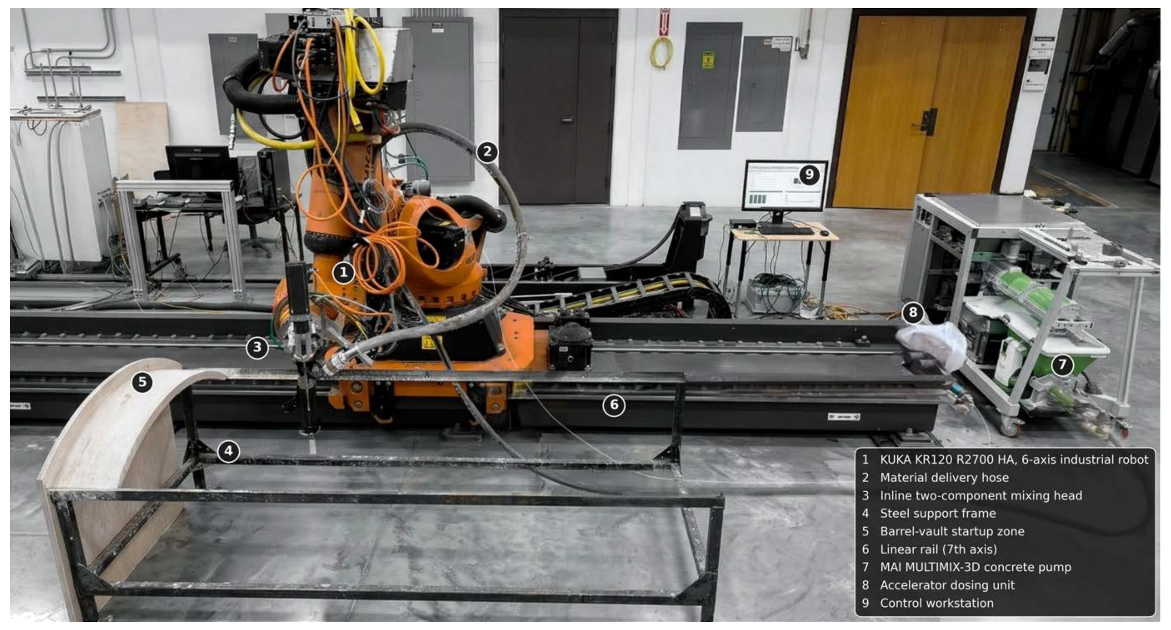

Figure 1: Fabrication setup. KUKA KR120 R2700 HA 6-axis robot with inline two-component mixing head, MAI MULTIMIX-3D concrete pump, and accelerator dosing unit. The 2K system enables programmable early-age kinetics — the process parameters predicted by the RL agent (accelerator dosage, flow rate, and TCP speed) map directly to these actuators.

As the print advances along the vault curvature, each new concrete filament is carried only by its interlayer bond to the one behind it, forming a growing cantilever. Once self-weight demand at the root interface exceeds the still-maturing bond capacity, the cantilever peels away and the print fails. Figure 2 shows this failure mode in a primary-only benchmark experiment: without lamination, the cantilevering filaments delaminate after only three filaments (42 mm of cantilever advance), demonstrating the fundamental cantilever limit of single-layer accelerated concrete.

A lamination layer — a secondary filament deposited in a zig-zag pattern over the primary layer — can prevent this failure by transferring load back into stable material. However, the lamination pattern, deposition schedule, and process parameters must be jointly determined so that stability is maintained at every location and time during fabrication. For a single vault this can be solved through expert tuning, but the solution is manual, geometry-specific, and does not transfer to new forms. The objective of this research is to learn a generalizable policy that automatically generates stability-feasible fabrication plans for previously unseen vaulted geometries.

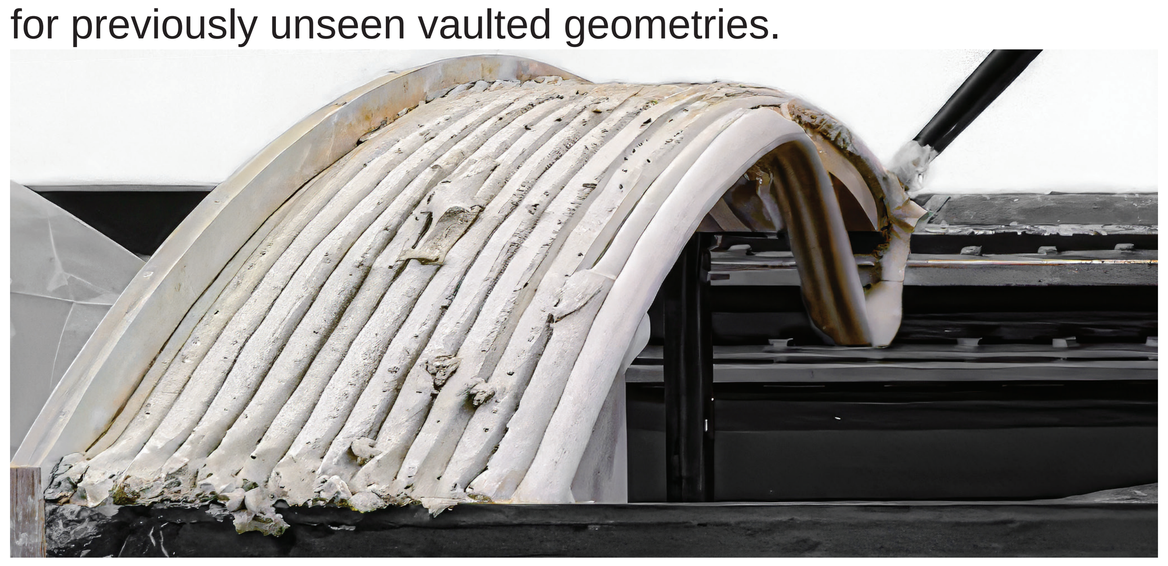

Figure 2: Primary-only benchmark (Experiment A). Without lamination, cantilevering filaments peel away from the root interface under self-weight after only 3 filaments (42 mm), demonstrating the fundamental cantilever limit of single-layer 3DCP.

02 Lamination Framework

The lamination design framework establishes the physics that the AI agent must learn to navigate. Construction proceeds in repeating cycles (Figure 3). First, primary filaments are deposited along the vault thrust line, extending the unsupported cantilever. A zig-zag lamination pass is then deposited over the primary layer. At this pre-bond state, the lamination is unbonded dead load on the cantilever — the root interface must survive the combined weight. Once bonded, the diagonal branches carry tension and redirect load back into the stable region, advancing the stable boundary. The cycle then repeats from the new stable boundary.

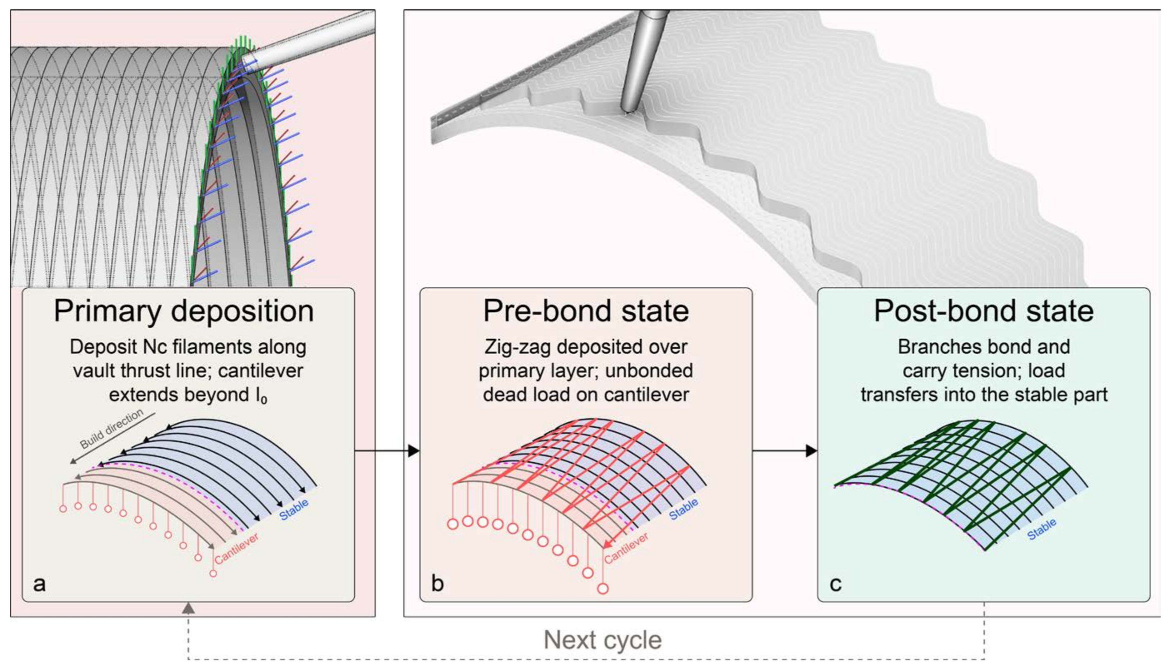

Figure 3: Lamination printing cycle. (a) Primary filaments are deposited along the vault thrust line, extending the unsupported cantilever. (b) A zig-zag lamination pass is deposited over the primary layer; at this pre-bond state the lamination is unbonded dead load. (c) Once bonded, diagonal branches carry tension and transfer load into the stable region, advancing the stable boundary. The cycle then repeats.

A zig-zag design is defined by two integers: the restraint-node count (how many attachment points are distributed along the filament length) and the tie-back depth (how far back into the stable region each branch reaches). These parameters create a central design trade-off (Figure 4): increasing the restraint-node count shortens each branch and steepens its inclination, improving the transverse force component but reducing bonded contact length. Increasing tie-back depth extends the transverse reach and lever arm but changes the dead-load distribution. Branch capacity transitions from bond-shear governed to material-strength governed at a crossover length, beyond which additional branch length adds only dead load with no capacity gain.

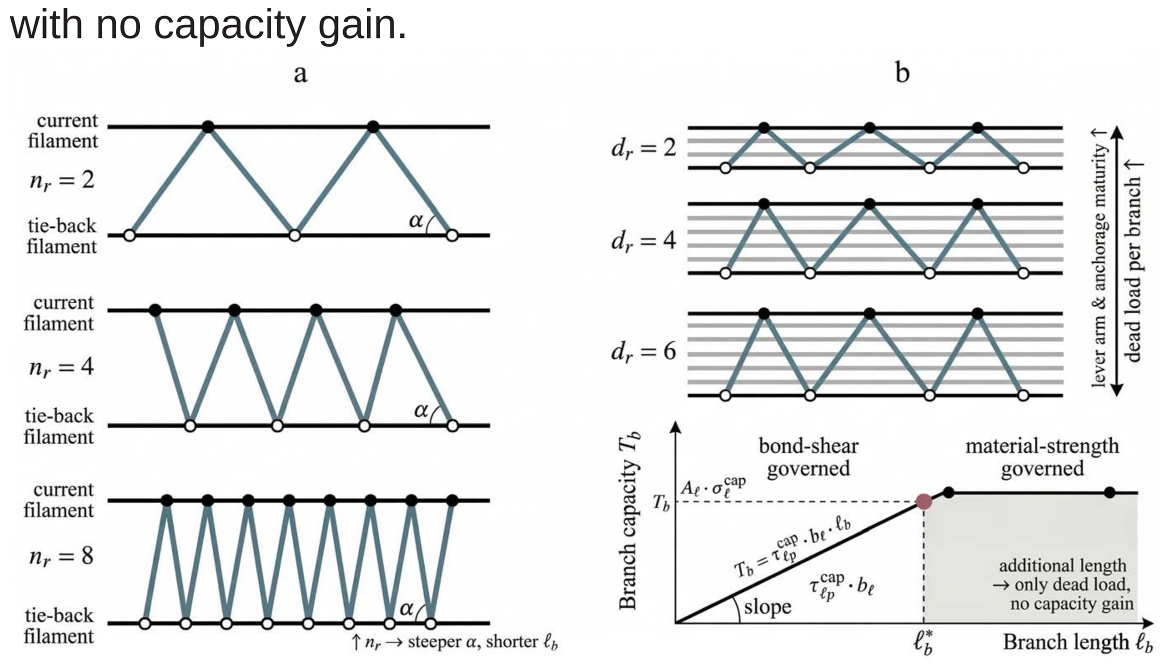

Figure 4: Zig-zag design parameters and trade-offs. (a) Effect of restraint-node count at fixed tie-back depth. (b) Effect of tie-back depth at fixed restraint-node count. (c) Branch capacity regimes: bond-shear governed versus material-strength governed, separated by the crossover branch length.

Admissibility requires passing two checks. Check 1 (deposition survivability) ensures the cantilever survives the pre-bond loading peak, when the freshly deposited zig-zag is unbonded dead load. Check 2 (chain sustainability) ensures the post-bond force balance prevents progressive collapse — each completed cycle must transfer enough force to compensate for the primary weight it introduces, otherwise a residual deficit accumulates across successive cycles and the chain of completed links eventually fails. The set of all designs passing both checks defines the admissible domain (Figure 5). The most robust design is selected by maximizing the minimum safety margin across both checks. Failed calibration tests fall outside the admissible region, confirming the predictive validity of the framework.

This admissible domain — parameterized by zig-zag restraint-node count, tie-back depth, and scheduling depth — becomes the constraint landscape that the RL agent must learn to navigate.

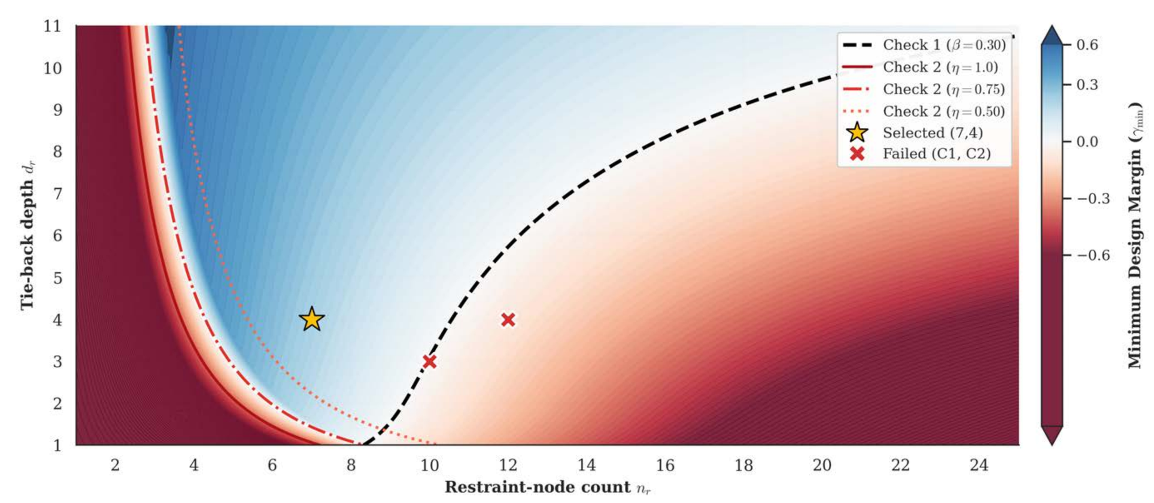

Figure 5: Admissible design domain. The blue region satisfies both Check 1 (pre-bond deposition survivability, dashed black boundary) and Check 2 (post-bond chain sustainability, red boundaries at three engagement fractions). The selected design (star) lies inside the admissible region; failed calibration tests (crosses) fall outside.

03 Experimental Validation

The lamination framework was validated through a full barrel-vault print using the selected design (Figure 6). The vault was completed scaffold-free, with primary filaments and zig-zag lamination passes alternating throughout the build. The top-view photograph shows the zig-zag pattern on the vault surface, confirming that the lamination branches engage across multiple underlying filaments as designed. This result demonstrates that the framework enables full vault completion for one specific geometry — the RL formulation aims to automate this design process for arbitrary vault forms.

Figure 6: Experiment B: scaffold-free barrel vault printed with the selected lamination design. (Top) Completed vault during printing, showing primary filaments and zig-zag lamination passes. (Bottom) Top view of the zig-zag pattern on the vault surface.

04 Why AI Is Necessary

The objective is not to find a stable print plan for one vault, but to learn a generalizable policy that automatically generates stability-feasible fabrication plans for previously unseen vaulted geometries. This requires jointly determining the lamination pattern, task schedule, and process parameters such that material capacity exceeds stress demand everywhere and at all times during construction.

Conventional geometric slicing derives toolpaths from form alone without considering time-dependent material behavior or stability constraints. Physics-based simulation can diagnose whether a given plan is stable but does not prescribe what action to take next. Neither method, in isolation, can produce the full set of coupled decisions required for automatic stability-aware planning across a family of freeform geometries.

Stability is inherently spatio-temporal: it depends on deposition history, the local age map, evolving capacity fields, and the chosen sequence of actions. Freeform vaults impose spatially non-uniform and continuously varying demand. Encoding this coupled dependence into fixed geometric rules produces brittle strategies that fail when the geometry changes. Reinforcement learning provides a systematic way to search the large decision space and learn transferable stabilization strategies.

05 RL Formulation

The AI system is formulated as a reinforcement learning problem in which an agent learns to make sequential deposition decisions by interacting with the spatio-temporal stability model (Figure 7). ***State.*** The agent observes the current deposited geometry, the spatial distribution of material age since deposition, and the corresponding capacity fields derived from calibrated material models. ***Action.*** At each decision step the agent selects which toolpath segment to deposit next, whether to continue the primary layer or switch to lamination, the zig-zag parameters (stitch count, spacing, amplitude), and the process parameters — accelerator dosage, flow rate, and TCP speed. ***Environment.*** The spatio-temporal stability model simulates the consequence of each action: it updates the geometry and age map, evaluates demand against time-dependent capacity, and returns the new state with a stability assessment. The admissibility checks from the lamination framework (Check 1 and Check 2) are encoded directly into the environment's evaluation logic. Reward. Positive reward is assigned when stability margins remain above threshold across all active interfaces. Negative reward is assigned for margin violations or simulated collapse.

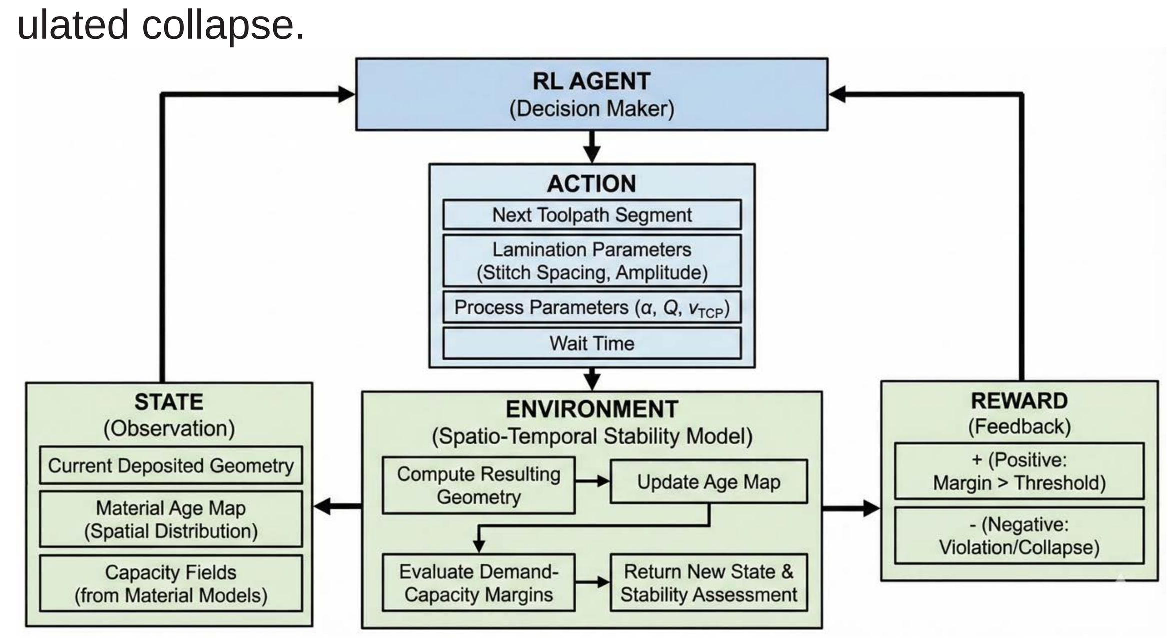

Figure 7: Reinforcement learning framework for stability-aware print-path generation. The agent selects toolpath segments, lamination parameters, process settings, and wait times. The spatio-temporal stability model acts as the environment, computing demand–capacity margins and returning a reward based on whether stability is maintained.