

Terrain-Adaptive Mobile 3D Printing with Hierarchical Control

Nors Shuangshan Li¹ and J. Nathan Kutz^{1,2}

¹Department of Electrical and Computer Engineering, University of Washington, USA

²Department of Applied Mathematics, University of Washington, USA

ssli93@uw.edu, kutz@uw.edu

Abstract -

Mobile 3D printing on unstructured terrain remains challenging due to the conflict between platform mobility and deposition precision. Existing gantry-based systems achieve high accuracy but lack mobility, while mobile platforms struggle to maintain print quality on uneven ground. We present a framework that tightly integrates AI-driven disturbance prediction with multi-modal sensor fusion and hierarchical hardware control, forming a closed-loop perception-learning-actuation system. The AI module learns terrain-to-perturbation mappings from IMU, vision, and depth sensors, enabling proactive compensation rather than reactive correction. This intelligence is embedded into a three-layer control architecture: path planning, predictive chassis-manipulator coordination, and precision hardware execution. Through outdoor experiments on terrain with slopes and surface irregularities, we demonstrate sub-centimeter printing accuracy while maintaining full platform mobility. This AI-hardware integration establishes a practical foundation for autonomous construction in unstructured environments.

Keywords -

Mobile 3D printing; AI-hardware integration; Hierarchical control; Terrain adaptation; Sensor fusion

1 Introduction

Automated construction has long been a human aspiration [1, 2]. It promises not only to accelerate labor-intensive tasks, but also to protect workers from danger by enabling construction in hazardous environments. Such capability is critical for establishing infrastructure in places that are otherwise inaccessible, including disaster recovery sites, remote locations, and challenging terrains. Existing robotic systems have begun to address parts of this challenge. Gantry-based systems, such as contour crafting and large-scale 3D printing [3, 4], have demonstrated the ability to fabricate buildings with high geometric precision. However, significant challenges remain in material properties, geometric conformity, and process robustness [5]. Similarly, fixed robotic arms deployed on-site [6] provide flexible control and accurate material deposition. However, both approaches face fundamental limitations:

gantry systems are immobile and require extensive setup, while fixed arms are restricted by their limited workspace. Both methods are therefore infrastructure-dependent and face challenges when applied to rapid deployment in unstructured environments.

Mobile robotic platforms offer a promising alternative by combining locomotion with manipulation capabilities [7, 8]. However, integrating high-fidelity construction with platform mobility remains a significant challenge [9]: ground irregularities induce vibrations and pose disturbances that compromise print quality, while the dynamic nature of mobile bases conflicts with the precision requirements of material deposition.

Here we present a hierarchical control framework that addresses these challenges by integrating data-driven learning with precision hardware control (Figure 1). The core idea is to leverage predictive models to learn environmental dynamics from sensor data and embed the learned knowledge into the hardware control loop, enabling intelligent real-time responses. Specifically, we developed a disturbance prediction module that learns terrain-to-perturbation mappings from multi-source sensors, predicting terrain-induced disturbances ahead of time to enable proactive compensation rather than reactive correction.


The main contributions of this work are:

- A three-layer hierarchical control architecture that separates path planning, predictive control, and hardware execution across appropriate time scales;
- Integration of data-driven disturbance prediction that enables proactive compensation for terrain-induced perturbations;
- Experimental validation demonstrating sub-centimeter printing accuracy on uneven terrain with full platform mobility.


2 System framework

By integrating data-driven models with real-time sensor feedback and precision hardware control, we developed a terrain-adaptive printing framework that simultaneously achieves both high mobility and high accuracy (Figure 2). The core philosophy is to enable the system to continuously adapt to environmental dynamics from sensor data


STEP 1: Adapting to uneven ground and printing the initial layer



STEP 2: Adjusting layer height to compensate for ground undulation



STEP 3: Leveling the base layer to ensure structural stability



STEP 4: High-speed printing with arm's real-time terrain adaption




Figure 1. Overview of terrain-adaptive mobile 3D printing. The mobile robot adapts to uneven ground and prints the initial layer (Step 1), adjusts layer height to compensate for terrain undulation (Step 2), levels the base layer to ensure structural stability (Step 3), and achieves high-speed printing with real-time terrain adaptation (Step 4).

and embed the learned predictive capabilities into the hardware control loop.

2.1 Data-driven disturbance prediction

The disturbance prediction module is the intelligent core of the framework. It learns terrain-to-perturbation mappings from multi-source sensors including IMU for chassis pose, RGB camera for end-effector tracking, and depth camera for terrain profiling. The module predicts 6-DOF disturbances $\hat{\mathbf{d}}_{t+\Delta t} \in \mathbb{R}^6$ (position and orientation perturbations) over a prediction horizon $\Delta t = 0.5$ s:

$$\hat{\mathbf{d}}_{t+\Delta t} = g(\mathbf{z}_t; \boldsymbol{\theta}) \quad (1)$$

where $\mathbf{z}_t \in \mathbb{R}^d$ represents fused sensor features (d ranges from 12 to 20 depending on sensor configuration) and $\boldsymbol{\theta}$ denotes learned model parameters. The input vector aggregates multi-modal information including IMU-derived linear acceleration and angular velocity, estimated base orientation, depth-derived terrain descriptors (local height variation and slope estimates), and end-effector pose relative to the mobile base. Raw images are not used directly; instead, compact geometric descriptors are extracted to enable real-time inference.

Sparse regression techniques [10] are applied for feature selection, identifying the most predictive sensor channels for disturbance estimation.

The disturbance prediction pipeline consists of two stages: a temporal encoding stage and a prediction stage.

The temporal encoder extracts latent representations of terrain-induced dynamics from sequential multi-modal sensor inputs. These encoded features are then processed by a lightweight feedforward network with two hidden layers of 128 and 64 units, which predicts future disturbances.

This design enables temporal modeling while maintaining low inference latency (approximately 12 ms), supporting real-time proactive compensation rather than reactive correction [11].

In Figure 2, the temporal encoding stage is represented by a SHallow REcurrent Decoders (SHRED) module [12, 13], implemented using Gated Recurrent Unit (GRU) blocks [14], while the downstream feedforward network corresponds to the disturbance prediction stage.

The shallow decoder (labeled as SDN in Figure 2) within the SHRED module maps recurrent latent states to compact feature representations.

The prediction module was trained on approximately 7 hours of operation data (4 hours from simulation, 3 hours from outdoor experiments) collected across four terrain types: flat concrete, grass-covered soil, loose gravel, and mixed terrain with slopes. Both real and simulated data were combined directly for training; the model's ability to generalize across these domains was validated through the outdoor experimental results. The dataset comprises approximately 1,400 trajectory segments totaling over 200 meters of traversed path, sampled at 50 Hz. Training and validation sets were split 80/20 at the trajectory level




Figure 2. Hierarchical control framework for terrain-adaptive mobile 3D printing. The system integrates adaptive path planning (left) with a perception-prediction-control loop (bottom). Sensor feedback from IMUs and vision-based tracking is processed through the disturbance prediction module, which enables feature selection for proactive compensation. The control policy generates real-time strategies for robotic actuation, enabling the mobile platform to maintain printing accuracy on unstructured terrain (right).

to avoid temporal correlation. Cross-time generalization was validated by resuming a printing task after a 24-hour interruption; the model maintained consistent prediction performance, demonstrating robustness to temporal discontinuities encountered in practical construction scenarios.

2.2 Three-layer hierarchical architecture

The prediction module is deeply integrated with the hierarchical hardware control architecture, forming a complete perception-learning-actuation closed loop:

Upper layer (0.1 Hz) handles adaptive path segmentation and velocity planning, transforming time-based printing trajectories into path-parameterized representations for

speed-independent precision control. The path parameterization maps the desired trajectory $\mathbf{p}(t)$ to a parameter-based form:

$$\mathbf{p}(s) = (1 - s)\mathbf{p}_0 + s\mathbf{p}_f, \quad s \in [0, 1] \quad (2)$$

where $\mathbf{p}_0$ and $\mathbf{p}_f$ denote initial and final positions, enabling velocity-independent trajectory tracking. Complex trajectories are decomposed into piecewise-linear segments at this layer. This layer dynamically adjusts path complexity based on geometric features and terrain conditions, ensuring smooth transitions between printing segments.

Middle layer (10 Hz) implements preview model predictive control (MPC), receiving disturbance predictions and coordinating chassis-manipulator motion [15, 16]. At

each control cycle, the MPC solves:

$$\min_{\mathbf{u}} \sum_{k=0}^N \left(\|\mathbf{e}_k\|_{\mathbf{Q}}^2 + \|\Delta \mathbf{u}_k\|_{\mathbf{R}}^2 \right) \quad (3)$$

subject to system dynamics $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k, \hat{\mathbf{d}}_k)$ and operational constraints (joint limits, velocity bounds), where $\mathbf{e}_k$ is tracking error, $\mathbf{u}_k$ is control input, and $\hat{\mathbf{d}}_k$ is the predicted disturbance from Eq. (1). The key innovation is a frequency-decomposition coordination strategy: the chassis handles low-frequency (<1 Hz), large-amplitude motions for terrain following, while the manipulator provides high-frequency (>5 Hz) precision corrections for accurate deposition.

Lower layer (100 Hz) provides precision hardware execution, controlling chassis steering, wheel speeds, and joint torques. High-frequency operation ensures real-time execution of MPC commands. Sensor data is fed back to the prediction module for real-time inference for disturbances adaptation.

In the current implementation, model weights for the disturbance prediction module are fixed during printing operations, and updates are performed intermittently using buffered operational data between printing sessions.

This design is chosen to ensure system stability and predictable control behavior during active material deposition, where continual online weight updates during printing could introduce transient prediction errors.

The system adapts to terrain conditions in real time through multi-modal sensor inputs, as the prediction module continuously processes multi-modal features to produce disturbance estimates without modifying model parameters.

This hierarchical architecture enables improved prediction performance across diverse terrains through offline model updates using accumulated operational data.

3 Validation

3.1 Simulation validation

To validate the hierarchical control framework before physical deployment, we conducted dynamics simulations in MATLAB/Simulink. The simulation environment modeled a 10-meter trajectory comprising a 5-meter flat section followed by a 5-meter slope at 5° inclination (Figure 3). To emulate terrain-induced disturbances, Gaussian noise with ± 5 mm amplitude was injected into the wheel transformation matrices (4×4 homogeneous transforms containing orientation and position). While real terrain disturbances exhibit spatially correlated, non-Gaussian characteristics, the simulation provides a controlled environment to verify the MPC's disturbance rejection capability; the outdoor experiments subsequently confirm generalization to realistic disturbance patterns.



Figure 3. Simulink simulation environment showing the mobile robot traversing a slope terrain. The trajectory includes a 5-meter flat section followed by a 5-meter slope at 5° inclination.

The MPC controller operated with a prediction horizon of 10 steps and control horizon of 5 steps, running at the middle control layer (10 Hz). Simulations were executed with a fixed-step solver at 1 ms resolution over a 30-second trajectory, representing continuous mobile printing operation.




Figure 4. Simulated end-effector position error over a 30-second trajectory. Despite continuous terrain disturbances from slope transition (dashed line) and injected noise, the MPC-based compensation maintains position errors below 5 mm in all three axes.

Figure 4 shows the end-effector position error throughout the simulated trajectory. Despite continuous terrain disturbances from both the slope transition and injected noise, the MPC-based compensation maintained position errors below 5 mm in all three axes. The error remained bounded without accumulation, validating the effectiveness of the predictive control strategy before field deployment.

3.2 Experimental setup

To further validate the terrain-adaptive capabilities on physical hardware, we conducted outdoor experiments on diverse terrain conditions. The experiments evaluated the system's ability to maintain printing precision despite ground irregularities.

The experimental platform consisted of a four-wheel drive mobile base (1.2 m $\times$ 0.8 m, 150 kg payload capacity) equipped with a 6-DOF robotic arm (1.5 m reach, 10 kg payload), as shown in Figure 5. Tests were conducted on terrain featuring slopes, surface roughness, and mixed surface types including grass, gravel, and concrete with irregularities. These conditions represent realistic challenges encountered in disaster recovery sites and undeveloped construction areas.



Figure 5. Experimental mobile 3D printing platform. Left: initial prototype with material hopper and extrusion system. Right: integrated system with 6-DOF robotic arm for terrain-adaptive printing.

Sensing systems include: IMU (100 Hz) for chassis state estimation, joint encoders (1000 Hz) and force/torque sensors (500 Hz) for manipulator control, stereo vision (30 Hz) for end-effector tracking, and RGB-D camera (30 Hz) for terrain mapping.

In addition, the printing process is conducted using a fixed nozzle diameter of 5 mm, a layer height of 4 mm, and a controlled nozzle travel speed of approximately 2cm/s.

These parameters are selected to maintain stable and consistent material deposition conditions throughout the experiments. As a result, variations in fabrication accuracy primarily reflect terrain-induced motion disturbances and the effectiveness of the proposed control framework, rather than variability in the material process.

This work focuses on kinematic disturbance compensation under uneven terrain. Material-related effects such as flow dynamics and rheology are assumed to be consistent and are not explicitly modeled in this study.

3.3 Performance results

In the primary validation experiment, the mobile robotic platform constructed a concrete foundation structure on uneven ground. Although limited in physical size, this structure represents a fundamental construction primitive that can be repeatedly deployed to form larger foundations or continuous wall segments—in practice, large-scale construction can be decomposed into such locally adaptive printing units. Despite challenging ground conditions, the hierarchical control framework maintained fabrication ac-

curacy within sub-centimeter level across the entire build, as confirmed by LiDAR scanning.

The disturbance prediction module proved critical for maintaining accuracy. By learning terrain-to-perturbation mappings, the system achieved effective disturbance rejection across different terrain types.

In this work, fabrication accuracy refers to the geometric deviation between the desired deposition trajectory and the actual realized deposition path.

Formally, it is defined as the Euclidean distance between the reference trajectory and the realized end-effector position:

$$\mathbf{e}(t) = p_{actual}(t) - p_{reference}(t) \quad (4)$$

The magnitude of the tracking error is given by its Euclidean norm:

$$\|\mathbf{e}(t)\|_2 \quad (5)$$

The deposition position is approximated by the end-effector position, assuming the extrusion offset was calibrated at the start of each printing session.

This definition focuses on geometric tracking accuracy under terrain-induced disturbances. Material-related effects such as flow dynamics, settling, and curing shrinkage are not considered in this metric, as they are independent of the control objective addressed in this work.

Table 1 summarizes end-effector tracking accuracy across different terrain conditions.

To clarify the reported metrics, all quantities in Table 1 are derived from the tracking error.

Specifically:

- *Height Dev.* denotes the mean and standard deviation of the vertical component of the tracking error (i.e., projection onto the z-axis).
- *Misalignment* represents the angular deviation between the desired and actual end-effector orientations.
- *Max Dev.* corresponds to the maximum Euclidean tracking error $\max_t \|\mathbf{e}(t)\|$ observed along the trajectory.

Table 1. End-effector tracking accuracy during printing on different terrain types

Terrain	Height Dev. (mm)	Misalign. (°)	Max Dev. (mm)
Flat concrete	2.1 $\pm$ 0.4	0.08 $\pm$ 0.02	3.8
Grass (soft)	4.3 $\pm$ 1.2	0.21 $\pm$ 0.08	8.7
Gravel	5.1 $\pm$ 1.5	0.25 $\pm$ 0.09	9.2
Slope (5°)	3.8 $\pm$ 0.9	0.18 $\pm$ 0.06	6.5
Mixed	4.7 $\pm$ 1.3	0.22 $\pm$ 0.07	8.9

Table 2 presents additional system performance metrics, demonstrating real-time capability and rapid disturbance recovery.

Table 2. System performance metrics

Metric	Value
MPC computation time	8.3 ± 2.1 ms
Prediction latency	12.5 ± 3.2 ms
Settling time	0.35 ± 0.12 s
Layer height consistency	± 0.8 mm

The final structure exhibited consistent layer stacking and mechanical stability. Across all terrain types, the system maintained sub-centimeter accuracy (mean position error < 5.1 mm), validating that the proposed framework enables mobile platforms to achieve precision approaching stationary gantry systems while operating on challenging terrain.

4 Discussion and conclusion

This study presents a hierarchical control framework that addresses the fundamental challenge of maintaining deposition precision on mobile platforms operating over unstructured terrain. Unlike existing approaches that either sacrifice mobility for accuracy or accuracy for mobility, our three-layer architecture achieves both through temporal separation of control concerns and predictive disturbance compensation.

The key technical contributions include: (1) a three-layer hierarchical control architecture that separates path planning, predictive control, and hardware execution across appropriate time scales; (2) integration of data-driven disturbance prediction that enables proactive compensation for terrain-induced perturbations; and (3) dual-layer chassis-manipulator coordination through frequency decomposition.

4.1 Failure modes and mitigation

Mobile construction systems operating on unstructured terrain face several characteristic failure modes. Understanding these failure modes and how the proposed framework addresses them is essential for assessing deployment readiness.

Terrain perception degradation occurs when environmental conditions compromise sensor accuracy. Varying lighting affects RGB-based end-effector tracking, while surface reflectivity and dust interfere with depth measurements. The multi-modal sensor fusion approach mitigates this by combining IMU, vision, and depth data, allowing the prediction module to maintain reasonable estimates even when individual sensors degrade. The learned terrain-to-perturbation mappings provide robustness by encoding statistical regularities rather than relying on instantaneous measurements alone.

Unlike laboratory robotic systems where disturbances

are isolated events, construction sites exhibit persistent environmental uncertainty: airborne dust from excavation activities, rapidly changing shadows from moving equipment, and surface glare from wet concrete or reflective materials. These conditions demand sensor fusion strategies that degrade gracefully rather than fail catastrophically. The hierarchical architecture addresses this by allowing the prediction module to operate with reduced feature sets when specific sensors become unreliable, falling back to IMU-dominated estimation during severe visual degradation while maintaining acceptable—though reduced—compensation performance.

Wheel slip and chassis compliance introduce discrepancies between commanded and actual platform motion, particularly on soft or loose surfaces. The frequency-decomposition coordination strategy addresses this by assigning terrain-following to the low-frequency chassis control loop while reserving the high-frequency manipulator for precision corrections. This separation ensures that chassis-level disturbances do not directly propagate to the deposition point.

Long-horizon drift poses a risk for extended construction tasks where small errors could accumulate over time. The experimental results provide direct evidence against this concern: over a 5-hour continuous printing session comprising approximately 140 layers, no systematic error accumulation was observed. This stability arises from the closed-loop architecture, in which sensor feedback continuously drives the prediction module’s real-time disturbance estimates, preventing drift from compounding across layers.

Material rheology and extrusion dynamics can further influence deposition quality; however, this work focuses on terrain-induced kinematic disturbances, and material-process interactions will be addressed in future studies.

4.2 Scalability implications

The dominant barriers to scaling mobile construction systems are typically long-duration stability and site-to-site variability, rather than geometric size alone. The proposed framework addresses both concerns through its hierarchical design.

For duration scaling, the demonstrated 5-hour operation without error accumulation suggests that the architecture handles the primary failure modes that typically prevent mobile systems from sustaining construction-time-scale operations. The locally adaptive nature of the control framework—where each printing segment is treated as a path-parameterized unit with real-time disturbance compensation—means that scaling to larger structures involves repeating proven primitives rather than extrapolating to untested behaviors.

For multi-robot scaling, the modular architecture pro-

vides a natural extension path. Each robot can operate its own perception-prediction-control loop while coordinating at the path-planning level. Recent work on aerial multi-robot construction [17] and collaborative mobile printing [18] demonstrates the feasibility of such coordination, and the hierarchical separation in our framework aligns well with distributed control paradigms.

From a practical deployment perspective, the system requires a minimum sensor configuration of one IMU and either stereo vision or an RGB-D camera for basic operation; the full sensor suite enhances performance but is not strictly necessary for all terrain types. The framework assumes locally continuous terrain variations with spatial frequencies below approximately 0.5 cycles per meter—corresponding to undulations with wavelengths greater than 2 meters. Extremely discontinuous obstacles such as steps, curbs, or debris piles exceeding the manipulator’s compensation range (approximately 15 cm vertical displacement) remain outside the current operational envelope and require path replanning at the upper control layer.

The construction workflow integration follows a modular primitive approach: rather than attempting continuous printing of entire structures, the system treats each foundation segment, wall section, or structural element as an independent printing unit with well-defined start and end conditions. This decomposition aligns naturally with construction sequencing practices, where curing time, material logistics, and inspection checkpoints create natural task boundaries. The demonstrated 5-hour operation capability corresponds to typical work shifts, suggesting compatibility with existing site management practices.

Several limitations remain before widespread deployment. The current system depends on accurate terrain mapping for optimal disturbance prediction, which may be challenging in rapidly changing environments such as active construction sites with concurrent human activity. Additionally, extending the framework to different construction materials will require characterizing material-specific disturbance responses. A systematic comparison with reactive-only control would further quantify the predictive module’s contribution.

Looking ahead, the practical implications of this work extend to disaster recovery, remote infrastructure construction, and eventually to extraterrestrial construction where autonomous operation on irregular surfaces is essential. By demonstrating that precision and mobility can coexist in construction robotics, this framework establishes a foundation for autonomous building in environments previously inaccessible to conventional methods.

References

- [1] Ke You, Cheng Zhou, Lieyun Ding, and Yuxiang Wang. Construction robotics in extreme environments: From earth to space. *Engineering*, 2025.
- [2] Robert Bogue. What are the prospects for robots in the construction industry? *Industrial Robot: An International Journal*, 45(1):1–6, 2017.
- [3] Sungwoo Lim, Richard A Buswell, Thanh T Le, Simon A Austin, Alistair GF Gibb, and Tony Thorpe. Developments in construction-scale additive manufacturing processes. *Automation in construction*, 21: 262–268, 2012.
- [4] Clément Gosselin, Romain Duballet, Ph Roux, Nadja Gaudillière, Justin Dirrenberger, and Ph Morel. Large-scale 3d printing of ultra-high performance concrete—a new processing route for architects and builders. *Materials & Design*, 100:102–109, 2016.
- [5] Richard A Buswell, WR Leal De Silva, Scott Z Jones, and Justin Dirrenberger. 3d printing using concrete extrusion: A roadmap for research. *Cement and concrete research*, 112:37–49, 2018.
- [6] Markus Giftthaler, Timothy Sandy, Kathrin Dörfler, Ian Brooks, Mark Buckingham, Gonzalo Rey, Matthias Kohler, Fabio Gramazio, and Jonas Buchli. Mobile robotic fabrication at 1: 1 scale: the in situ fabricator: system, experiences and current developments. *Construction Robotics*, 1(1):3–14, 2017.
- [7] Steven J Keating, Julian C Leland, Levi Cai, and Neri Oxman. Toward site-specific and self-sufficient robotic fabrication on architectural scales. *Science robotics*, 2(5):eaam8986, 2017.
- [8] Abdullah Alhijaily, Abdulrahman Alqarni, Zekai Murat Kilic, and Paulo Bartolo. Development of a mobile 3d printer and comparative evaluation against traditional gantry systems. *Journal of Intelligent Manufacturing*, 36(6):3783–3800, 2025.
- [9] Kathrin Dörfler, Gido Dielemans, Stefan Leutenegger, Selen Ercan Jenny, Johannes Pankert, Julius Sustarevas, Lukas Lachmayer, Annika Raatz, and Dirk Lowke. Advancing construction in existing contexts: Prospects and barriers of 3d printing with mobile robots for building maintenance and repair. *Cement and Concrete Research*, 186:107656, 2024.
- [10] Eurika Kaiser, J Nathan Kutz, and Steven L Brunton. Sparse identification of nonlinear dynamics for model predictive control in the low-data limit.

Proceedings of the Royal Society A, 474(2219): 20180335, 2018.

- [11] João M Silva, Gabriel Wagner, Rafael Silva, António Morais, João Ribeiro, Sacha Mould, Bruno Figueiredo, João M Nóbrega, and Paulo JS Cruz. Real-time precision in 3d concrete printing: controlling layer morphology via machine vision and learning algorithms. *Inventions*, 9(4):80, 2024.
- [12] Jan P Williams, Olivia Zahn, and J Nathan Kutz. Sensing with shallow recurrent decoder networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 480(2298), 2024.
- [13] Mars Liyao Gao, Jan P Williams, and J Nathan Kutz. Sparse identification of nonlinear dynamics and koopman operators with shallow recurrent decoder networks. *arXiv preprint arXiv:2501.13329*, 2025.
- [14] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [15] Sotaro Katayama, Masaki Murooka, and Yuichi Tazaki. Model predictive control of legged and humanoid robots: models and algorithms. *Advanced Robotics*, 37(5):298–315, 2023.
- [16] Marios-Nektarios Stamatopoulos, Jakub Haluška, Elias Small, Jude Marroush, Avijit Banerjee, and George Nikolakopoulos. Fully autonomous chunk-based aerial additive manufacturing with offset-free predictive control. *Automation in Construction*, 178: 106361, 2025.
- [17] Ketao Zhang, Pisak Chermprayong, Feng Xiao, Dimos Tzoumanikas, Barrie Dams, Sebastian Kay, Basaran Bahadir Kocer, Alec Burns, Lachlan Orr, Talib Alhinai, et al. Aerial additive manufacturing with multiple autonomous robots. *Nature*, 609 (7928):709–717, 2022.
- [18] Shuai Li, Tian Lan, Hung-Xuan Nguyen, and Phuong Tran. Frontiers in construction 3d printing: self-monitoring, multi-robot, drone-assisted processes. *Progress in Additive Manufacturing*, 10(4):2001–2030, 2025.