AI-Based Process Planning for Robotic Fabrication of Customized Orthopedic Implants Under Material and Geometric Constraints

Nataniel Gwardiak
Department of Computer Science, University of New Hampshire, Durham, NH, USA.
natanielgwardiak@unh.edu

Makayla Malone
Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
contactmakayla@colostate.edu

Abstract

Customized orthopedic implants must satisfy patient-specific geometry, stiffness, porosity, and manufacturability requirements that are difficult to reconcile in conventional sequential CAD/CAM workflows. This paper develops an AI-based process-planning framework for robotic fabrication of such implants under coupled material and geometric constraints. The implant volume is represented as a patient-specific graph of local geometry, lattice morphology, and material descriptors; a learned surrogate ensemble predicts defect risk, fabrication time, and support demand for candidate plans; and a constrained optimizer selects tool poses, deposition or machining parameters, and path sequencing subject to reachability, collision, minimum feature size, porosity, and effective modulus limits. The formulation treats manufacturability as a risk-averse constraint rather than a post hoc check, so geometric accuracy and mechanical compatibility are evaluated jointly with robot kinematics. Analytical treatment of the resulting feasible set clarifies why rule-based planning fails when porosity, strut thickness, and tool orientation interact, and it identifies the conditions under which surrogate error changes the selected plan. The framework yields a defensible decision rule for balancing shape fidelity, fabrication cost, and defect risk, and it provides a technical basis for integrating patient imaging, material models, and robot programming into a single planning loop. The main contribution is a constraint-aware architecture and mathematical decision model, not a validated clinical workflow; its significance lies in making the trade-offs among implant function, material behavior, and robotic execution explicit and computable before fabrication.

Keywords

robotic fabrication; orthopedic implants; process planning; geometric constraints; material constraints; constraint-aware machine learning

1. Introduction

Additive and robotic fabrication have expanded the design space for customized orthopedic implants by allowing biocompatible metals and other materials to be deposited, consolidated, or machined in patient-specific forms [1]. A persistent clinical concern is stress shielding: implants that are much stiffer than surrounding bone can alter load transfer and reduce bone density near the interface [2]. Patient-specific finite-element studies have shown that implant geometry and material distribution can be adjusted to reduce such mismatch, but

the resulting shapes are often difficult to translate into reliable fabrication plans [3]. The planning problem is therefore not merely a geometric slicing exercise; it requires simultaneous reasoning about anatomical shape, mechanical function, material architecture, and the physical reach and motion of a robot.

Porous and lattice architectures are central to this problem because they can lower effective stiffness while preserving strength and providing space for bone ingrowth. Cellular-solid theory relates the effective modulus of an open-cell structure to its relative density and the modulus of the solid strut material [4]. This relationship is useful for implant design, but it also creates a coupling that process planning cannot ignore: the porosity that improves mechanical compatibility also changes minimum feature size, heat transfer, deposition stability, and the risk of local defects. A plan that satisfies nominal geometry may produce struts that are too thin for the selected nozzle or beam, or may create local thermal histories that alter the intended material state.

Machine learning has been applied in manufacturing to predict process outcomes, detect anomalies, and support decision-making from high-dimensional sensor and design data [5]. Robotic fabrication platforms add another layer of complexity because the end effector, workpiece orientation, and tool path are programmable in six or more degrees of freedom, enabling non-planar deposition and multi-axis access [6]. These capabilities are relevant to customized implants, where accessible surfaces are curved and support structures may be difficult to remove. However, most planning pipelines still separate geometric design, material selection, toolpath generation, and robot programming. That separation makes it difficult to evaluate whether a locally attractive porosity distribution is actually reachable, collidable, and stable during fabrication.

The central gap addressed here is the absence of a planning formulation that treats material and geometric constraints as joint conditions on robot-executable plans. Robot motion planning provides mature methods for reachability and collision avoidance [7]. Surrogate modeling provides a way to approximate expensive fabrication or physics-based evaluations during optimization [8]. Mesh optimization for spatial wireframe robotic 3D printing further shows that print paths can be treated as a constrained graph rather than a fixed output of slicing, which is directly relevant to non-planar and wireframe-like implant architectures [9]. Building on these ideas, the present work formulates process planning as a constrained decision problem in which a learned model estimates local manufacturability risk and an optimizer selects a plan that is feasible under material, geometric, and kinematic limits.

Surrogate-based analysis and optimization offer a principled route to embedding expensive evaluations within design loops [10]. Uncertainty estimation is essential in this setting because a surrogate can be confidently wrong outside its training distribution; Bayesian approximations and ensemble disagreement provide practical signals for risk-averse decisions [11]. Material tailoring in electron beam melting has demonstrated that implant stiffness and architecture can be adjusted through process and design choices [12], while rapid-prototyping principles clarify how layer, path, and support decisions interact with geometric fidelity [13]. The contribution of this paper is a constraint-aware AI process-planning framework and a mathematical decision model for robotic fabrication of customized orthopedic implants. The framework integrates patient-specific geometric representation, material constraints, robot feasibility, learned risk estimation, and constrained optimization. The analysis focuses on the structure of the feasible set, the conditions under which surrogate

error changes the selected plan, and the technical trade-offs that determine whether a plan is both fabricable and clinically meaningful.

2. Problem Formulation and Constraints

The implant is represented by a target volume Ω with boundary ∂Ω, a patient-specific bone interface region, and a desired local material distribution. A fabrication plan π specifies the sequence of tool poses, path segments, deposition or machining parameters, and support strategy. The plan is executable only if it satisfies three families of constraints: geometric accuracy relative to the target shape and interface, material constraints on porosity and effective mechanical behavior, and kinematic constraints imposed by the robot, tool, and fixture. The feasible set is written as

$$\Pi_feas = \{\pi \in \Pi \mid C_geo(\pi) \leq 0, C_mat(\pi) \leq 0, C_kin(\pi) \leq 0\}$$

where Π is the set of candidate plans, C_geo aggregates shape deviation and minimum feature-size violations, C_mat aggregates porosity, modulus, and material-continuity constraints, and C_kin aggregates reachability, joint-limit, and collision conditions. The inequality form is convenient because each constraint can be evaluated separately and then combined in a constrained optimizer. The set Π_feas may be empty; in that case the design or fixture must be revised rather than forcing an infeasible plan.

Material constraints are particularly important for orthopedic implants because stiffness mismatch can dominate clinical performance. For an open-cell lattice with solid modulus E_s, porosity φ, and empirical exponent n, a standard cellular-solid approximation gives an effective modulus E_s(1 − φ)^n. The implant planning problem can require this modulus to lie within a window that is low enough to reduce stress shielding but high enough to preserve structural integrity:

$$E_min \leq E_s(1-\varphi)^n \leq E_max$$

Here E_min and E_max define the acceptable effective modulus range, and n depends on cell topology and deformation mode. The relationship is an established approximation for open-cell foams, not a patient-specific prediction; it is used to expose the coupling between porosity and mechanical behavior. If the planner chooses a porosity that satisfies this inequality but creates struts below the minimum printable or machinable thickness, the geometric constraint C_geo will reject the plan. This interaction is a primary reason why sequential design-then-slicing workflows can produce nominally valid implants that are not fabricable.

Robot feasibility introduces a second geometric layer. For each tool pose p_i in a plan with N poses, the tool center point must maintain a safety distance δ from obstacles O, including the fixture, the patient model, and the robot itself, while the joint vector q_i must remain within joint limits Q. A compact statement of the kinematic condition is

where d(·,·) is the minimum Euclidean distance between the tool geometry at pose p_i and the obstacle set, and Q denotes the admissible joint space. This condition is necessary but not sufficient for fabrication: collision-free poses may still produce unacceptable orientation changes, excessive acceleration, or inaccessible support removal. Motion-planning methods provide the computational basis for evaluating such conditions [14]. In the proposed framework, the kinematic constraint is evaluated inside the planning loop rather than after toolpath generation, so candidate material distributions that are unreachable can be rejected before detailed process simulation.

3. Methodological Framework

3.1 Patient-specific geometric and material representation

The framework begins with a patient-specific volume derived from imaging and implant design data. The volume is discretized into a graph whose nodes represent local material cells, lattice junctions, or surface patches, and whose edges encode geometric adjacency and material continuity. Node features include local wall thickness, distance to the bone interface, surface curvature, lattice coordination number, and material phase. Edge features include path direction, connectivity, and estimated support demand. This graph representation is more suitable than a fixed voxel grid for robotic fabrication because it preserves the connectivity of struts and toolpath segments, which determines whether a plan can be executed as a continuous motion.

The graph is not a substitute for geometric computation. It is an interface between patient-specific design and the planning model. Geometric constraints such as minimum feature size, interface tolerance, and accessible surface area are computed on the original boundary representation and then mapped to graph nodes. Material constraints are evaluated on local descriptors such as porosity, strut diameter, and estimated effective modulus. The mapping must be conservative: if a node is near a constraint boundary, the planner should treat it as uncertain rather than assign it a precise nominal value.

3.2 Material constraints and mechanical compatibility

For metallic implants, laser and electron-beam powder-bed processes can produce controlled porosity and tailored architectures, but the achievable range depends on powder size, beam energy, scan strategy, and post-processing [15]. The planner therefore needs a material-feasibility model that distinguishes between a desired porosity and a fabricable porosity. Small struts may be unresolved by the beam or may lack sufficient mechanical integrity; large struts may increase stiffness and reduce the benefit of the porous architecture. These effects are local, but they propagate through the lattice connectivity. A node-level violation can create a disconnected load path or a stress concentration that invalidates the intended effective modulus.

The material model used in planning should be conservative in two respects. First, it should use lower bounds on achievable strut thickness and upper bounds on local porosity gradients. Second, it should propagate uncertainty from process parameters to effective properties rather than treating the nominal lattice as exact. This is consistent with the use of surrogate models in engineering design, where expensive simulations are approximated by models that are fast enough for optimization but must be accompanied by error estimates [8], [10]. The proposed framework does not assume that a single material model is universally valid; it allows alternative constitutive approximations to be plugged into the same constraint structure.

3.3 Geometric constraints and toolpath feasibility

Geometric constraints include shape deviation from the target boundary, minimum feature size, interface fit, and support accessibility. In robotic fabrication, the toolpath is a sequence of poses and process parameters, not merely a set of curves. The orientation of the end effector affects collision clearance, deposition direction, and the ability to reach concave regions. A path that is geometrically short may require large orientation changes that exceed joint velocity limits or produce unstable deposition. Conversely, a path that is kinematically

smooth may deviate from the desired surface. The planner must therefore evaluate geometric fidelity and robot motion together.

This coupling is especially relevant to customized implants because the bone interface is curved and the accessible workspace is constrained by the fixture and patient anatomy. A plan that is feasible for a flat test coupon may be infeasible when the implant is mounted in a patient-specific orientation. The graph representation allows the planner to encode accessibility as a node-level feature and to penalize paths that require excessive reorientation. The penalty is not arbitrary; it reflects the physical cost of accelerating and decelerating the robot and the increased risk of collision near the fixture.

3.4 AI-based planning architecture

The planning architecture combines a learned surrogate ensemble with a constrained optimizer. The surrogate receives a descriptor z for a local plan segment and predicts a defect risk r_ψ(z), where ψ denotes the learned parameters. The risk is interpreted as the probability or relative likelihood that the segment violates a local material or geometric requirement, such as insufficient fusion, excessive porosity, or unacceptable shape deviation. A logistic output layer is a convenient choice because it maps a learned score g_ψ(z) to the unit interval:

$$r_\psi(z) = \sigma\big(g_\psi(z)\big)$$

where σ(·) is the logistic function. The score g_ψ(z) can be produced by a graph neural network or another differentiable model that consumes node and edge features. The surrogate is not treated as ground truth. It is used to rank candidate plans and to identify regions where further physics-based evaluation is needed.

To quantify model uncertainty, the framework uses an ensemble of K surrogate models with parameters ψ_k. The ensemble mean and variance for a local segment are

$$\mu(z) = \left(\frac{1}{K}\right)\sum_{k=1}^{K} r_{\psi_k}(z)$$

$$\sigma^2(z) = \left(\frac{1}{K}\right)\sum_{k=1}^{K} \big(r_{\psi_k}(z) - \mu(z)\big)^2$$

where μ(z) is the mean predicted risk and σ²(z) is the ensemble variance. High variance indicates that the models disagree, which is a practical warning that the segment lies near a decision boundary or outside the region well represented by training data. The variance is not a formal posterior probability, but it provides a useful risk signal in the same spirit as Bayesian approximation methods for deep learning [11]. The planner can then impose a risk-averse constraint of the form

$$\mu(z) + \lambda\sigma(z) \le \varepsilon$$

where λ ≥ 0 controls aversion to uncertain risk and ε is the maximum acceptable risk for a local segment. When λ is zero, the constraint uses only the mean prediction. As λ increases, the planner rejects segments whose predicted risk is acceptable on average but highly uncertain. This is essential for customized implants because a small number of local defects can compromise mechanical integrity or create surfaces that are difficult to clean and sterilize.

The optimizer selects a plan by minimizing a composite objective subject to the feasible-set and risk-averse constraints. A representative scalarization is

$$\min_{\pi} J(\pi) = w1T(\pi) + w2V_s(\pi) + w3D(\pi) + w4R(\pi)$$

where T(π) is estimated fabrication time, V_s(π) is support or sacrificial material volume, D(π) is a shape-deviation measure, and R(π) is an aggregated risk penalty. The weights w1 through w4 are nonnegative and sum to one. The choice of weights is a design decision that should reflect clinical priorities: a plan that minimizes fabrication time may be unacceptable if it increases interface deviation, while a plan that minimizes deviation may require supports that are difficult to remove from a porous implant. The objective is evaluated only over plans that satisfy the kinematic constraints and the material-feasibility conditions. This structure prevents the optimizer from trading an infeasible collision clearance for a small improvement in shape fidelity.

The architecture also incorporates process-planning knowledge rather than treating the robot as a generic motion device. The integration of design, process, and execution decisions is consistent with open and intelligent CNC process-planning frameworks, which emphasize interoperable decision layers rather than isolated post-processors [18]. In the present framework, the learned surrogate does not replace physics-based simulation; it reduces the number of expensive evaluations needed to search the plan space. Candidate plans that pass the surrogate screen are then checked against the full constraint set. This two-stage logic is important because surrogate error is nonuniform: predictions are often more reliable in dense regions of the training distribution and less reliable near constraint boundaries.

Material deposition or machining parameters are included in the plan vector. For deposition processes, the volumetric flow rate can be approximated as the product of path speed, layer thickness, and bead width. This relationship is not a process model, but it exposes the coupling between speed, geometry, and material supply. If the planner increases speed to reduce fabrication time, it must also adjust flow or thickness to maintain continuity; otherwise the local risk of voids or excess material increases. Robotic deposition of soft and structured materials illustrates how tool motion and material rheology interact, and why a purely geometric path can fail when the material cannot support the intended shape [17]. The proposed framework accommodates such effects through local descriptors and risk constraints rather than through a single global process model.

4. Analytical Assessment and Technical Discussion

The feasible set defined by the joint constraints is generally nonconvex. Porosity appears in the material constraint, minimum feature size appears in the geometric constraint, and tool orientation appears in the kinematic constraint. Increasing porosity can improve modulus compatibility but reduce strut thickness below the geometric limit. Reducing strut thickness can improve bone ingrowth potential but increase defect risk. Reorienting the workpiece can improve reachability but change support requirements and thermal boundary conditions. These interactions explain why rule-based planning, which applies constraints in a fixed sequence, can miss feasible plans or accept plans that fail under a different constraint ordering. A plan that satisfies the geometric constraint after slicing may violate the material constraint after thermal distortion; a plan that satisfies the material constraint may be unreachable by the robot.

The risk-averse constraint $\mu(z) + \lambda\sigma(z) \leq \varepsilon$ has a direct interpretation. For a fixed ε, increasing λ shrinks the set of admissible local segments. The planner becomes more conservative in regions where the surrogate ensemble disagrees. This is desirable when the cost of a local defect is high, as in load-bearing regions or at the bone interface. It is less

desirable in noncritical regions where excessive conservatism may eliminate plans that are actually fabricable. A practical implementation could vary λ by anatomical region, using a larger value near the interface and a smaller value in sacrificial support regions. The framework allows this because the risk constraint is local and can be aggregated with region-specific weights.

Surrogate error changes the selected plan when two candidate plans are close in objective value but differ in predicted risk. If the risk difference is small relative to ensemble variance, the ranking is unstable. The optimizer can reduce this instability by preferring plans with lower variance, by requesting additional physics-based evaluation for high-variance segments, or by selecting a plan that is robust across the ensemble. The last option is equivalent to minimizing an upper confidence bound on risk, which is a standard risk-averse decision rule. This does not eliminate error, but it makes the decision sensitive to uncertainty rather than relying on a single point prediction.

Material and geometric constraints also impose a trade-off between support volume and shape fidelity. Supports are often necessary to prevent distortion in overhanging regions, but they can be difficult to remove from porous implants and can leave surface defects. A plan that reduces support volume by reorienting the implant may increase collision risk or require non-planar paths that are harder to execute. The objective function makes this trade-off explicit through $V_s(\pi)$ and $D(\pi)$. The weights should not be treated as universal constants; they encode clinical and manufacturing priorities. For example, interface regions may prioritize $D(\pi)$, while noncritical surfaces may tolerate greater deviation in exchange for lower support volume and shorter fabrication time.

The kinematic constraint $\min_i d(p_i, O) \geq \delta$ is conservative because it evaluates distance at discrete poses. Continuous collision checking would be more accurate but computationally expensive. A practical compromise is to evaluate discrete poses during surrogate screening and then perform continuous checking on the selected plan. This preserves the speed of the AI-based search while reducing the risk of missed collisions between poses. Similarly, material constraints can be evaluated at graph nodes and then refined with finite-element or process simulation in regions flagged by high ensemble variance. The framework is therefore best understood as a decision architecture that allocates computational effort where it matters, not as a replacement for validated physical models.

The clinical relevance of the framework lies in making trade-offs explicit. Customized implants are often designed under competing requirements: the surgeon wants an intimate fit, the material scientist wants controlled porosity, and the fabricator wants a stable process. A process plan that satisfies one requirement while silently violating another is a poor plan, even if each subsystem reports success. The proposed formulation forces the planner to state the acceptable modulus window, the minimum feature size, the collision margin, and the risk tolerance before optimization. These are not purely technical parameters; they reflect clinical and manufacturing policy. By embedding them in a single constrained decision model, the framework provides a basis for auditing why a particular plan was selected and which constraints were active.

Several boundary conditions limit the claims that can be made. The cellular-solid approximation is valid only for open-cell structures within a certain porosity range and does not capture anisotropy, defects, or post-processing effects. The surrogate ensemble requires representative training data; outside that distribution, variance may increase but may not fully reflect model error. The kinematic model assumes a rigid robot and a static fixture, whereas

real fabrication involves thermal distortion, tool deflection, and fixture compliance. The objective weights are normative choices, not empirical constants. These limitations do not undermine the planning formulation, but they define the conditions under which its outputs should be interpreted. The framework is most defensible when used as a constraint-aware decision aid that identifies feasible regions and exposes trade-offs, rather than as an autonomous clinical decision-maker.

5. Conclusion

This paper has developed an AI-based process-planning framework for robotic fabrication of customized orthopedic implants under material and geometric constraints. The central contribution is a formulation that treats the fabrication plan as a constrained decision variable rather than a downstream product of design. Patient-specific geometry is represented as a graph of local material and geometric descriptors; a surrogate ensemble estimates local defect risk and uncertainty; and a constrained optimizer selects tool poses, process parameters, and path sequencing while enforcing porosity, effective modulus, minimum feature size, reachability, and collision constraints. The risk-averse constraint $\mu(z) + \lambda\sigma(z) \leq \varepsilon$ provides a principled way to reject plans whose predicted performance is acceptable on average but unreliable in local regions.

The analysis shows that the feasible set is shaped by interactions among porosity, strut thickness, tool orientation, and support demand. Rule-based planning can fail because these interactions are order-dependent and because a locally valid decision may violate a global kinematic or material condition. The proposed architecture addresses this by integrating learned risk estimation with explicit constraint evaluation. The surrogate reduces the cost of searching candidate plans, while the constraint model preserves physical meaning and prevents the optimizer from exploiting prediction error. The objective function makes trade-offs among fabrication time, support volume, shape deviation, and risk explicit, allowing clinical and manufacturing priorities to be encoded rather than assumed.

The framework does not claim a validated clinical workflow. Its significance is methodological: it provides a computable structure for reasoning about implant function, material behavior, and robotic execution before fabrication begins. By making constraints and uncertainty explicit, the approach supports traceable planning decisions and identifies where additional physical modeling or process data would most affect the selected plan. For customized orthopedic implants, where geometric fit and mechanical compatibility are both essential, such constraint-aware planning is a necessary step toward reliable robotic fabrication.

References

1. Murr, L. E., Gaytan, S. M., Ramirez, D. A., Martinez, E., Hernandez, J., Amato, K. N., Shindo, P. W., Medina, F. R., & Wicker, R. B. (2012). Metal fabrication by additive manufacturing using laser and electron beam melting technologies. Journal of Materials Science & Technology, 28(1), 1–14.

2. Niinomi, M., & Nakai, M. (2011). Titanium-based biomaterials for preventing stress shielding between implant devices and bone. International Journal of Biomaterials, 2011, 836587.

3. Harrysson, O. L. A., Hosni, Y. A., & Nayfeh, J. F. (2007). Custom-designed orthopedic implants evaluated using finite element analysis of patient-specific computed

tomography data: Femoral-component case study. BMC Musculoskeletal Disorders, 8, 91.

4. Gibson, L. J., & Ashby, M. F. (1997). Cellular solids: Structure and properties (2nd ed.). Cambridge University Press.

5. Wang, J., Ma, Y., Zhang, L., Gao, R. X., & Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. Journal of Manufacturing Systems, 48, 144–156.

6. Keating, S. J., & Oxman, N. (2013). Compound fabrication: A multi-functional robotic platform for digital design and fabrication. Robotics and Computer-Integrated Manufacturing, 29(6), 439–448.

7. Siciliano, B., Sciavicco, L., Villani, L., & Oriolo, G. (2010). Robotics: Modelling, planning and control. Springer.

8. Forrester, A. I. J., Sóbester, A., & Keane, A. J. (2008). Engineering design via surrogate modelling: A practical guide. Wiley.

9. Yu, L., & Luo, D. (2018). Mesh optimisation for spatial wireframe robotic 3d printing. International Journal of Rapid Manufacturing, 7(2-3), 169-185.

10. Queipo, N. V., Haftka, R. T., Shyy, W., Goel, T., Vaidyanathan, R., & Tucker, P. K. (2005). Surrogate-based analysis and optimization. Progress in Aerospace Sciences, 41(1), 1–28.

11. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In Proceedings of the 33rd International Conference on Machine Learning (Vol. 48, pp. 1050–1059). PMLR.

12. Harrysson, O. L. A., Cansizoglu, O., Marcellin-Little, D. J., Cormier, D. R., & West, H. A. (2008). Direct metal fabrication of titanium implants with tailored materials and mechanical properties using electron beam melting technology. Materials Science and Engineering: C, 28(3), 366–373.

13. Chua, C. K., Leong, K. F., & Lim, C. S. (2010). Rapid prototyping: Principles and applications (3rd ed.). World Scientific.

14. Latombe, J.-C. (1991). Robot motion planning. Springer.

15. Sing, S. L., An, J., Yeong, W. Y., & Wiria, F. E. (2016). Laser and electron-beam powder-bed additive manufacturing of metallic implants: A review on processes, materials and designs. Journal of Orthopaedic Research, 34(3), 369–385.

16. Goh, G. D., Sing, S. L., & Yeong, W. Y. (2021). A review on machine learning in 3D printing: Applications, potential, and challenges. Artificial Intelligence Review, 54, 63–94.

17. Truby, R. L., & Lewis, J. A. (2016). Printing soft matter in three dimensions. Nature, 540(7633), 371–378.

18. Xu, X., & Newman, S. T. (2006). Making CNC machine tools more open, interoperable and intelligent—A review of the technologies. Computers in Industry, 57(2), 141–152.