

AgenticDfAM: An agentic workflow automates design for additive manufacturing through traceable adaptive assessment and multicriteria redesigns

David Guirguis

d.g.ca@ieee.org

Carnegie Mellon University <https://orcid.org/0000-0001-5341-2437>

Research Article

Keywords: Large-Language-Models, Design for additive manufacturing, Agentic artificial intelligence, Geometric reasoning, Computer-aided design, Process-informed redesign

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11434386/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

AgenticDfAM: An agentic workflow automates design for additive manufacturing through traceable adaptive assessment and multicriteria redesigns

David Guirguis

Mechanical Engineering Department, Southern Methodist University, Dallas, TX
dguirguis@smu.edu

Abstract

Design for Additive Manufacturing tools apply predefined geometric checks, while dimensional limits require experimental validation for specific alloys, machines, and process settings. Engineers often interpret their applicability and manually coordinate assessment, redesign, and simulation. Large language models alone cannot reliably provide precise geometric measurements. This paper introduces AgenticDfAM, which coordinates reasoning agents and tools to revise designs and their assessment and construction procedures. Tools retain authority over numerical evidence. The feature-recognition and measurement tools were developed and benchmarked against reference geometries. Scripted agent replies were used to test validation of communication between agentic stages, error handling, and recovery. Invalid inter-agent messages were rejected by semantic validators. Across computational case studies, three redesign strategies remove all reference violations while exposing material and distortion tradeoffs. Thermal feedback lowers a part’s predicted layer-scale peak temperature by 160 °C. In a support design experiment, support above an intended contact surface was identified by a reviewer and corrected by the agent after diagnosis of the construction code. The correction was verified on curved and flat targets and reused in four autonomous clamp runs, with returns to support construction reduced from ten to five. The workflow investigates issues beyond prescribed checks while retaining verifiable engineering evidence.

Keywords: Large-Language-Models, Design for additive manufacturing, Agentic artificial intelligence, Geometric reasoning, Computer-aided design, Process-informed redesign

Highlights

- AgenticDfAM adapts designs and the procedures used to assess and construct them
- Geometric measurements ground agent reasoning in verifiable numerical evidence
- Manufacturing guidance is translated into executable, rechecked design changes
- Process feedback guides redesign across material, thermal, and distortion tradeoffs
- Agent-directed investigations extend assessment beyond the initial rule catalog

1 Introduction

Design for additive manufacturing (DfAM) is used to relate manufacturing requirements to part geometry, build orientation and support [1–4]. In metal additive manufacturing (AM), limits on wall thickness, overhang angle and unsupported span are established for particular alloys, machines and process settings [3, 4]. Their applicability must therefore be checked before they are transferred to another manufacturing setup [4]. Assessment requires both a geometric measurement and an interpretation of the experimental conditions supporting its limit. When new guidance becomes available, the rule set and the procedures used to apply it must also be updated.

Substantial steps of DfAM assessment and redesign have been automated. Geometry is assessed against manufacturing guidance through geometric screening, feature recognition and knowledge-based systems [5–8]. Support demand has been reduced by mesh deformation [9], horizontal holes have been converted into teardrops [10], parametric surfaces have been optimized under AM constraints [11], and predicted distortion has been compensated through geometric changes [12]. Part geometry, build orientation and support layout have also been optimized together [13], with transient thermoelastic models used for support design [14]. In these approaches, the assessment and redesign are defined by the selected rules, objectives and constraints (Fig. 1a).

When separate design, assessment and simulation tools are combined, applicable guidance must be selected, geometry and results transferred, and redesigns reassessed. These tasks are often coordinated by engineers [1, 5, 10, 12]. The engineering judgment extends beyond a pass or fail: removal of a geometric violation can increase material use or change thermal and mechanical response [1, 13, 14]. An adaptive DfAM workflow must therefore connect measurements to their applicability and to the consequences of a proposed change. When the available checks do not resolve a concern, further analysis or revision of the assessment or construction procedure may be required.

Geometrical features and manufacturability have been predicted using learning models (Fig. 1b) [15–18], and learned labels have been combined with measured computer-aided design (CAD) attributes in hybrid approaches [19]. Technical guidance and visual geometry can also be interpreted using large language models, although errors in rule retrieval, dimensional interpretation, and feature recognition have been documented [20, 21]. These findings motivate explicit verification of language-model judgments. Precise geometric measurements and numerical comparisons should be obtained through executable tools, with their inputs and results retained for review. A model can propose a dimension or identify a concern, while the corresponding measurement must establish the numerical evidence.

Agentic workflows have been developed in which reasoning, tool use and revision are interleaved in response to observations [22] (Fig. 1c). For DfAM, this provides a route to connect guidance interpretation, geometric measurement, executable redesign and process feedback within an engineering task. The question addressed here is whether these capabilities can be combined so that manufacturing assessments and design alternatives are revised through agent reasoning grounded in accurate geometric and computational evidence. Both the design and the procedures used to assess or construct it can then be adapted, while the numerical basis of each engineering decision remains open to verification.

This paper introduces AgenticDfAM, an adaptive engineering workflow in which accurate geometric measurement is coupled with agent reasoning to direct manufacturing assessment and multicriteria redesign informed by published manufacturing guidance and process simulations (Fig. 2). Guidance is interpreted by agents, geometric quantities and rule comparisons are computed by developed tools, and proposed changes are executed and rechecked on CAD solids. Assessment, redesign and process simulation are connected by stage review, with further analyses or procedural corrections requested when the returned evidence requires them. Submissions are validated for record structure, rule and geometry references, measurement coverage and permitted redesign parameters. Geometric compliance, material use, predicted maximum temperature and build distortion are considered together when alternatives are evaluated. Reasoning can therefore direct changes to geometry and to the procedures used to assess and construct it, with the supporting measurements and computational results retained for inspection.






Fig. 1. Conventional DfAM tools, machine learning models and AgenticDfAM. (a) Geometric checks are applied against selected criteria, and build orientation and support structures are prepared or optimized using specified settings [5, 13]. (b) CAD face labels are predicted by a supervised model [16], while assessments are generated from images and requirements by a multimodal large language model (LLM) [20]. The dashed arrow indicates training supervision. (c) Guidance, the rule catalog and run records are interpreted by LLM agents to select tool requests and review returned evidence [22]. Geometric measurements, numerical checks, CAD operations and simulations are supplied by developed tools. Assessments and design alternatives are generated by agents, with returned evidence used to revise designs and the assessment or construction procedures. Orange arrows indicate requests and agent outputs. Blue arrows indicate returned evidence. Bidirectional arrows connect the agents with context and the tools with project files. CAD, computer-aided design.

a Knowledge and geometric assessment



b Design, simulation and feedback



c Reviewer Agent

Saved-design investigation




Fig. 2. AgenticDfAM workflow. (a) Manufacturing guidelines are converted into a rule catalog by the Rules Librarian. CAD features are recognized and measured using CAD tools. Feature identification and measurements are reviewed by the Feature Analyst, and the assessment against DfAM criteria is examined by the Compliance Reviewer. (b) Build orientation precedes geometric redesign or support construction. Both routes are supported by shared tools and stage-review interfaces. Thermal and mechanical simulations are used to inform further redesign and compare alternatives. Engineering assessments and design alternatives are generated by the agents using numerical evidence from the tools. (c) Stage outputs are reviewed against task requirements using evidence from developed tools. Unexpected results are investigated, and corrections are requested and verified.

2 Methods

2.1 Knowledge representation and agent roles

The IN718 guideline [4] was located and interpreted by the Rules Librarian. A typed catalog of 15 entries spanning ten check types was generated from the paper. Executable numerical limits were encoded in 13 entries, including the advisory demonstrator thickness. Flow-channel and lattice recommendations, C1 and L1, were retained as advisory text without executable limits. Applicable classes, quantity, units, comparison and strictness, severity, source location, supporting passage and tested envelope were stored for each rule. Each finding was assigned a violation, warning, pass, advisory or unevaluated status. For measured x and limit x_{lim} , the demand ratio was defined as

$$D = \max \left(\frac{x_{\text{lim}}}{x}, \frac{x}{x_{\text{lim}}} \right) \quad (1)$$

where a ratio was omitted when its limit was not specified. $D > 1$ therefore indicates a measurement beyond either limit. Each rule of the reference catalog specified one governing limit, so D reduced to one of the two ratios. The 80° value stored with A1 marks the steepest tested angle and was not applied as a limit.

Each verdict was obtained from its recorded inequality. A criterion was not met at $D > 1$ for a non-strict limit or $D \geq 1$ for a strict limit. D was reported for ranking within each rule and was not used to assign verdicts or compare manufacturing difficulty across rules. For a lower limit, $D = \infty$ was returned when $x \leq 0$. Flat downward-facing surfaces at $\alpha = 0^\circ$ were assessed by O1 or B1. Wall class was assigned from dimensions and end support. Advisory guidance was retained, and conditions requiring unavailable measurements were reported as unevaluated.

Four primary agent roles were assigned. The rule catalog was prepared by the Rules Librarian. Recognized features were examined by the Feature Analyst using geometric tools. Checker findings were assessed by the Compliance Reviewer, and redesign alternatives were proposed by the Design Reviser. Typed records were exchanged between these roles. Identifiers, required measurements, review coverage, parameter ranges and distinct executable alternatives were enforced by Pydantic schemas and semantic checks. Measurement, numerical comparison, CAD execution and rechecking were reserved for the tools. Applicability, explanations and strategy selection were assigned to the agents (Supplementary Fig. S9). Changes to applicability were permitted for documented feature reclassification, support context or a recognized design detail. Severity was recorded separately from the numerical criterion. The support-volume threshold of 25% was designated as a study warning for material use. An exceedance of an extrapolated limit could be retained as a warning with its applicability rationale. When agent judgments were combined with checker findings, the original values and demand ratios were preserved. Request limits of 8, 40, 40, and 40 were assigned to the Rules Librarian, Feature Analyst, Compliance Reviewer, and Design Reviser, respectively. Twelve requests were allowed for each stage reviewer. Three retries were permitted for stream-parsing failure.

The rule catalog and reviewed features were submitted to the rule checker, and its findings were assessed by the Compliance Reviewer before redesign alternatives were proposed. Orientation, redesign, support construction and simulation were assessed against stage catalogs. Only support-catalog review returned work during the complete clamp runs. The orientation, redesign and simulation catalogs were applied once after each run and reported their findings without returning work to those stages. For live investigation and saved-design review, geometric queries and rendered views were made available, and concerns outside the catalog could be submitted with applicability reasons and requested corrections (See Fig. 2). Measurements and numerical comparisons were computed by tools, and submissions were validated. Published limits, transferred parameters and study screening thresholds are distinguished in Supplementary Note 3.

2.2 Geometry and measurements

Geometric evidence for the agents was obtained from the CAD models using developed tools. STEP models were read as single solids using Open CASCADE through OCP [23]. Distances and surface parameters were queried from the boundary representation. For views, ray queries, and selected sampled measurements, a triangulation was generated with 0.05 mm linear and 0.25 rad angular deflection. For a unit build direction $\mathbf{b}$, vertex set V and point $\mathbf{p}$, the height relative to the part's lowest vertex was defined as

$$h(\mathbf{p}) = \mathbf{p} \cdot \mathbf{b} - \min_{\mathbf{v} \in V} (\mathbf{v} \cdot \mathbf{b}) \quad (2)$$

Cylindrical faces were grouped by common axis, radius, and convexity. A complete cylinder was identified at angular coverage of at least 359° . Horizontal and vertical features were distinguished by axis angles to the plate of at most 5° and at least 85° , respectively. Through and blind holes were distinguished using end probes. Tubes were identified from coaxial inner and outer cylinders, and teardrop holes from cylindrical surfaces with tangent planar roofs. Recognition and sampling tolerances are given in Supplementary Note 2.1.

Planar face pairs were retained when the dot product of their outward unit normals was below -0.99999 . For these antiparallel faces with centroids $\mathbf{c}_i$ and $\mathbf{c}_j$ and outward normal $\mathbf{n}_i$, the signed separation is defined as

$$s = (\mathbf{c}_j - \mathbf{c}_i) \cdot \mathbf{n}_i \quad (3)$$

A candidate wall thickness $t = -s$ was assigned for negative s , and a gap $g = s$ for positive s . Rays were required to pass between the two faces through material or free space, excluding pairs whose projected regions did not overlap. End probes were used to determine wall support. For a downward-facing surface, its angle α to the plate was defined as

$$\alpha = \arccos |\mathbf{n} \cdot \mathbf{b}| \quad (4)$$

where $\mathbf{n}$ is the unit surface normal. Downward-facing surfaces satisfy $\mathbf{n} \cdot \mathbf{b} < 0$, and a horizontal underside is assigned $\alpha = 0^\circ$. For a boundary edge e of a flat downskin, support by the neighboring face G was identified when

$$z_{\min}(G) < z(e) - \varepsilon \quad \text{and} \quad \mathbf{n}_G \cdot \mathbf{d}_{\text{in}}(e) > 0.5 \quad (5)$$

Here, height is measured along $\mathbf{b}$ as z , ε is set to 0.05 mm, and the unit direction into the underside across the edge is denoted $\mathbf{d}_{\text{in}}(e)$. This criterion identifies the edge from which unsupported reach is measured. A connection to the build plate was assessed separately through support anchoring. A bridge was identified from opposed supported edges, and its span was defined as

$$b_s = \min_{(i,j)} |(\mathbf{m}_j - \mathbf{m}_i) \cdot \mathbf{d}_i| \quad (6)$$

where edge midpoints are denoted $\mathbf{m}_i$ and $\mathbf{m}_j$ and the inward direction is denoted $\mathbf{d}_i$. Candidate edge pairs were retained when $\mathbf{d}_i \cdot \mathbf{d}_j < -0.99$. For other supported configurations, the governing cantilever reach was defined as

$$L = \max_{p \in S} \min_{e \in E} d(p, e) \quad (7)$$

where surface samples and vertices are contained in S , supported edges are contained in E , and the shortest point-to-edge distance is denoted $d(p, e)$. Faces were sampled at 1.5 mm spacing, reduced to 0.75 mm for faces smaller than 200 mm^2 , and every triangulation vertex was included. Multiple supporting edges were included in the reach calculation (Supplementary Fig. S6).

The STL bracket was stitched into a faceted solid. Coplanar faces were merged, and cylindrical regions were reconstructed before recognition. For remaining curved downskins, the angle of the

flattest substantial facet was used. Chamfers and fillets were classified separately and were excluded from the initial feature catalog because no rule was assigned to them. In the clamp study, downward-facing surfaces were also checked independently of feature labels, so geometry outside a rule-bearing feature could be examined.

2.3 Redesign and verification

Six initial operator families were implemented for teardrop-hole cuts, pitched bridge roofs, ribs, downskin filling or trimming, wall thickening and tube thickening (Supplementary Table S14). For a circular bore of radius r and tangent roof angle θ measured from the build plate, the apex rise above the original crown was defined as

$$\Delta h = r \left(\frac{1}{\cos \theta} - 1 \right) \quad (8)$$

The hole-roof angle was capped at 70° because the rise in Eq. (8) diverges as θ approaches 90° . For a triangular rib extending a horizontal reach R from a wall to a feature at height h , its inclined lower edge descended toward the wall at angle φ to the build plate. The rib-wall intersection height was defined as

$$h_{\text{attach}} = h - R \tan \varphi \quad (9)$$

Material below the lower rib-wall intersection was required in the supporting wall unless the rib reached the plate. For a plane steepened to θ across the vertical interval $z_{\text{top}} - z_{\text{bot}}$, the horizontal reach is defined as

$$a_\theta = \frac{z_{\text{top}} - z_{\text{bot}}}{\tan \theta} \quad (10)$$

Geometric feasibility was checked using these relations before Boolean execution. Clearances of 0.5 mm above roofs and 0.2 mm beside wall changes were required. A valid single solid was required after each operation. Before the next change, the geometry was recognized again and the target was relocated by type and position, because face identifiers could be changed by Boolean operations. Wall and tube changes were applied first, followed by ribs, downskins, bridge roofs and hole roofs. Each submitted alternative was executed in preview and rechecked with the original catalog.

Three strategies were generated by the analytical-artifact planner: local changes with principal-envelope retention in V1, additional treatment of warnings that could be addressed geometrically in V2, and material removal by trimming in V3. Each alternative was executed, recognized again and checked against the reference catalog. Simulation-informed bracket changes were generated by the agent in the coding-agent environment. Local stress and thermal concentrations were used to identify candidate changes, and common process settings were used to compare their geometric, material and simulated consequences. Functional equivalence was not tested.

Clamp supports were constructed from a target envelope, 0.2 mm grid walls and contact teeth, with part material subtracted. An initial grid hatching of 1.625 mm was adopted from the reported block-support study [24]. Tooth height, top length, base length and spacing were set to 0.4, 0.7, 1.0 and 0.3 mm following the tooth parameters of IN625 as the closest available material resource to IN718 [25]. Support-grid spans under SUP-G1 were sampled at 0.15 mm during construction and on the exported design. Permitted parameters were revised in the support loop using measured unsupported spans. Additional computation or operator revision could be requested for a concern outside the catalog. In the demonstrated correction, each surface triangle selected as requiring support was swept from its actual position to the plate. The effects were verified with a flat-target control, a cylindrical-target comparison, complete directional reruns and review of the saved geometry (Supplementary Notes 8.2 and 8.3).

2.4 Numerical process simulations

Layerwise build simulations were performed with Ansys Additive and Mechanical APDL 2025 R2. IN718 was simulated at 280 W laser power, 1000 mm/s scan speed, 95 μm hatch distance, 60 μm deposited-layer thickness, and 20 $^{\circ}\text{C}$ plate temperature [4]. Documented software defaults were used for beam diameter and stripe settings. Elastic-plastic mechanics with bilinear isotropic hardening were driven by the vendor machine-learning thermal-strain model. A generic machine definition was used for the analytical artifact because its approximately 315 mm envelope exceeded the build area of the guideline machine. Mechanical voxel sizes were set to 1.5 mm for the analytical artifact, 0.5 mm for the bracket and 1 mm for the clamp. The clamp support wall thickness was one-fifth of the 1 mm voxel size, so clamp displacement fields were treated as exploratory predictions. The support correction was assessed separately through CAD measurements (Supplementary Note 8.2). Fields were read from solver results and mapped to model coordinates. Maximum displacement and the 95th percentile of absolute surface-normal deviation were reported among the dimensional measures.

For the transient thermal analysis, temperature-dependent material properties, an absorptivity of 0.5, and a convection coefficient of $12.5 \text{ W}/(\text{m}^2\text{K})$ were used. Absorptivities of 0.3 and 0.7 were also evaluated [26]. Thermal concentrations were identified using a relative peak-temperature criterion adapted from prior work [27, 28], together with a section-reduction criterion specific to this implementation. The normalization and geometric criteria are specified in Supplementary Note 7. These temperatures were interpreted as part-scale heat accumulation. Surface-normal deviation was obtained by projecting surface displacements onto each design’s undeformed CAD normals. Model settings, mechanical fields, dimensional measures and thermal sensitivity calculations are given in Supplementary Notes 6.2, 6.3, 6.4 and 7.

2.5 Evaluation and execution provenance

Four geometries were evaluated: a purpose-designed analytical artifact, CMU bracket [29], NIST artifact [30] and a clamp model obtained from GrabCAD [31]. The IN718 reference catalog was applied to all four. Analytical-artifact measurements were compared with a separate STEP-entity parser. The NIST reference key was prepared from published dimensions and STEP parameters with language-model assistance and author review, separately from the recognizer. After the initial comparison, two entries were corrected for a rounded base area and a missing wall limit. No verdict was changed by either correction. Measurement agreement was assessed at resolutions of 0.001 mm, 0.001° or 0.001 mm^2 . Measurement agreement and feature context were assessed in the geometric comparisons. In the adaptive support experiment, one defect was followed through investigation, diagnosis, code correction, a flat-target regression control, and four complete directional reruns. Reuse of the same correction was assessed in those runs.

The rule checker and redesign planner were used for the early geometric benchmarks, catalog comparisons and initial artifact redesigns. For protocol controls, scripted model replies were substituted while the orchestrator, tools, validation and combination of findings were retained.

During the support investigation, the defect was identified, geometric queries were selected, the construction code was diagnosed and the operator was corrected by the agent without author guidance. The correction was locally verified before four autonomous runs were performed, from STEP input to simulated alternatives, with no human intervention during execution. Support-catalog review was applied during the runs. The orientation, redesign and simulation catalogs were applied once after each run for reporting. Language-model review was subsequently applied to the saved support designs by the Support Reviewer, using geometric queries, rendered views and documented applicability reasons. Measurement, numerical comparison, CAD execution and submission validation were supplied by developed tools (Supplementary Notes 5.3, 8.2 and 8.3). Claude Opus 5 was used for the computational tasks identified above. The implementation was executed on Windows 11 with an Intel Core Ultra 9 285K processor, 192 GB memory and an NVIDIA GeForce RTX 5090. Python [32], OCP [33], VTK [34], NumPy [35], Matplotlib [36], and Pydantic [37] were used.

3 Results

3.1 Validation of feature recognition and geometric measurements

The feature-recognition and measurement tools were evaluated on the National Institute of Standards and Technology (NIST) additive manufacturing test artifact [30] (Fig. 3). A reference key was prepared from the published geometry with language-model assistance and author review, separately from the recognizer. Seventy of the 71 features specified in the NIST drawing and included in the benchmark were recognized. Eighteen features present in the STEP geometry but not separately specified in the drawing were also included in the reference key. These comprised 17 walls, gaps and thin skins in the first fine-feature set and one contact region between the base and build plate. All 18 were recognized. Agreement with the reference values was obtained for all 226 quantities at the reporting resolution. However, the central through bore was not recognized as a separate feature. The central through bore, 10 mm in diameter and 17 mm long, was not assigned a separate feature record because its cylindrical surface was paired with a coaxial outer surface as a 3 mm-high tube. Six walls were also classified differently in terms of end support. No catalog verdict was changed by these limitations, because no binding vertical-hole rule was assigned to the bore and the same thickness limit was applied to the walls.

To compare the feature recognizer against an existing learned method, the NIST artifact was processed with BRepNet [16], a neural network that operates directly on the boundary representation (B-rep) of a CAD model, the set of faces, edges and vertices that a STEP file stores. BRepNet passes information between neighbouring faces and edges and, in the pretrained model released with the Fusion 360 Gallery segmentation dataset, assigns each face one of eight modeling-operation classes, such as extrude side, cut end, fillet or chamfer. These labels describe how a face was created in the CAD history, not what it means for printing. BRepNet does not group faces into DfAM features such as holes, thin walls or overhangs, and it has no notion of build direction, so its output could not be scored with the feature-level accuracy used for our recognizer, and no common accuracy score is reported (Supplementary Note 1.4). The reference key for the artifact was compiled within this study from the published NIST drawing.

A purpose-designed analytical artifact was used to examine support conditions as well as measurement agreement. Holes, thin walls, slots, a tube and inclined and horizontal overhangs were included in its 40 analytic faces. Eighteen features were recognized, and all 24 dimensional and angular comparisons agreed with values obtained using a separately implemented STEP-entity parser at resolutions of 0.001 mm and 0.001°. With the positive CAD Y axis used as the build direction, 38 findings were produced, including 11 violations and five warnings. The horizontal cantilever was supported along two adjacent edges. Its governing unsupported reach was measured as 40 mm to the nearest supporting edge, whereas a 50 mm projection would have been obtained from one edge alone. Support context was thus required to identify the quantity to which the limit was applied (Supplementary Notes 1.2 and 2.2).

Faceted geometry was also assessed using CMU bracket [29] that consists of a base, a central web with four openings and a cylindrical boss (Supplementary Fig. S13). Fifty-three features were recognized after reconstruction of its 5,038 triangles. Under the reference catalog, no violation and one roughness warning were obtained. Fourteen regions were flagged by a facet screen using the same limits. Of these, 13 were associated with chamfers, blends or walls accepted or excluded under the feature-based policy. The difference was attributed to feature context and rule assignment. Whether the identified features would cause a build failure was not established (Supplementary Note 1.3).

Orientation was also found to affect the analytical-artifact assessment. Eleven violations were recorded at a build height of 110 mm in +Y, compared with three violations at 70 mm in either +Z or -Z. Features were assessed relative to the plate in each direction, so compliance could not be assigned to the geometry independently of orientation (Supplementary Note 1.5).




Fig. 3. Feature recognition on the NIST additive manufacturing artifact [30]. (a) Published geometry and recognized feature groups. (b) Reference entries by group and assessment status. The central through hole was not recognized separately because its bore surfaces were grouped with a coaxial tube. (c) Lateral features and their rule findings. (d) Fine-feature sets of slots, bosses, pins and holes. The reference correspondence and recognition limitations are documented in Supplementary Note 1.1.

3.2 Sensitivity of manufacturing judgments to guideline interpretation

DfAM guidance must be interpreted against the manufacturing process and the requirements of the intended product. A successful build alone does not establish compliance with dimensional, surface-finish or performance requirements. In the IN718 guideline [4], inclined walls were fully built down to 25°, although substantially greater downskin roughness was measured at smaller inclination angles. Surface-finish advice may therefore remain a recommendation for one application but become binding when incorporated into the requirements of another. Requirements may be imposed through contracts following AWS D20.1/D20.1M [38] and MIL-STD-31000C [39]. For applicable spaceflight AM processes, NASA-STD-6030 requires surface texture and detail resolution acceptance criteria to be defined as part of material process qualification [40]. Sensitivity to guideline interpretation was examined by holding geometry and measured quantities fixed while changing how recommendations and experimental limits were encoded. The resulting differences indicate which manufacturing judgments depend on those interpretations.

Manufacturing judgments were examined separately from geometric measurement. An IN718 laser powder bed fusion guideline [4] was located and interpreted by the Rules Librarian. Fifteen

catalog entries spanning ten check types were generated. The source passage, applicable feature classes, units, comparison convention, and tested conditions were retained for each entry. Measured quantities were compared with the encoded criteria, with findings assigned violation, warning, pass, advisory, or unevaluated status. The interpretation connecting a published statement to an executable criterion was therefore available for inspection (Supplementary Note 3). Small walls 15 mm high and 10 mm long and large walls 30 mm high and 25 mm long were tested in the guideline [4]. Minimum thicknesses of 0.30 and 0.60 mm were specified for these two classes, respectively.

The effect of interpretation was assessed by applying four documented policies to unchanged geometry and measured values (see Fig. 4). Selected recommendations and advisory limits were made binding under the strict policy and relaxed under the permissive policy. Under the envelope policy, the large-wall thickness limit was applied as a violation to walls longer than the small test sample, rather than as a warning under the reference policy. The original analytical artifact, CMU bracket and NIST artifact were included, together with three analytical-artifact redesigns: V1, envelope retention, V2, warning treatment, and V3, material removal. On the original artifact, 6 to 17 violations were obtained, compared with 11 under the reference policy. On the three redesigns, all reference violations were removed, but 3 to 9 violations were obtained under the strict policy. The differences followed from the treatment of recommendations, experimental limits and dimensions beyond the tested specimens. Compliance was therefore conditional on the documented interpretation as well as the measured geometry. These alternatives were prescribed policies, rather than repeated language-model extractions (Supplementary Note 4).




Fig. 4. Sensitivity to guideline interpretation on unchanged geometry. Four documented policies were applied to the same geometry and measurements. Selected recommendations were made binding under the strict policy and relaxed under the permissive policy, and the large-wall thickness limit was applied as a violation under the envelope policy. (a) Violations for each part. The black tick marks the count under the reference policy, and arrows show the count under the strict, permissive, and envelope policies. An open circle marks a policy that gives the same count as the reference. The grey band marks compliance. (b) Warnings, shown in the same way. (c) Feature verdicts changed against the reference policy, by feature kind, summed over the six parts. Verdicts made stricter are drawn to the right and relaxed verdicts to the left. Inclined surfaces accounted for 20 of the 38 changes under the strict policy and 20 of the 28 under the permissive policy. The analytical-artifact redesigns are V1, envelope retention, V2, warning treatment, and V3, material removal. Policy definitions are given in Supplementary Note 4.

3.3 Validation of inter-agent communication

The use of verified measurements in an agentic workflow also depends on the validity of inter-agent communication. Typed records were exchanged through the Rules Librarian, Feature Analyst, Compliance Reviewer and Design Reviser interfaces. Record structure was enforced by schemas, while rule and geometry references, required measurements and proposed changes were checked for semantic consistency. Invalid messages and attempts to alter numerical evidence were tested separately.

In a protocol test, eight deliberately invalid calls or submissions were included among 17 requests and 20 tool calls (Fig. 5). Each invalid item was rejected, and the sequence was completed after corrected responses were supplied. Rule definitions, feature and face references, review coverage, redesign parameters and retention of a targeted violation were tested. When a Compliance Reviewer request was declined in a separate control, the checker findings were returned with 11 violations and five warnings intact. Validation, feedback, and recovery were demonstrated with scripted replies.

A bridge with a measured span of 10 mm and a 3 mm limit was used to test whether agent messages could override a violation identified by the checker. A pass without an admissible applicability reason was submitted in one message, and an altered span was submitted in another. Both messages were rejected. After an admissible review was combined with the checker findings, all 38 numerical values and demand ratios were preserved. On the 40-face analytical artifact, recognition was completed in 348 ± 3 ms and rule checking in 0.26 ± 0.01 ms over ten repetitions. These timings covered tool execution, with model latency excluded.




Fig. 5. Scripted test of inter-agent communication and tool execution times. (a) Requests made through the four agent interfaces in the scripted protocol test on the analytical artifact, in order from left to right. Seventeen requests and 20 tool calls were made, and eight deliberately invalid calls or submissions were included. Each box is coloured by what was returned to the agent. Numbered boxes contain an injected fault, numbered as in Supplementary Table S13. Each fault was returned as an error, and the sequence was completed after corrected responses were supplied. (b) Execution times for the tools and report generation, given as mean $\pm$ standard deviation next to each point. Ten runs were used, except for report generation, for which three runs were used. Model latency is excluded.

3.4 Geometric redesign and material tradeoffs

Three redesign strategies were applied to the analytical artifact using the developed planner, with targets set by the geometric criteria derived from [4]. The design alternatives are shown in Fig. 6. Six operator families were available for hole roofs, bridge roofs, ribs, downskins, walls and tubes. Local changes were made while the principal envelope was retained in V1. Additional warnings that could be addressed geometrically were treated in V2. Material was removed from selected shallow overhangs in V3. Eleven, 13 and 11 operations were applied, respectively. Valid single solids were obtained, and all 11 violations were removed on re-recognition and rechecking, with nine, one and nine warnings retained. Volume was increased by 28.0% in V1 and 54.0% in V2, but reduced by 15.2% in V3. Compliance with the same reference rules was therefore obtained with different material and envelope consequences (Supplementary Note 6.1).

3.5 Dimensional consequences of redesign

Different dimensional outcomes were predicted for the three redesigns with no reference-catalog violations (Fig. 7). Maximum displacement after plate release was reduced from 4.84 mm for the original to 3.02, 2.01 and 2.67 mm for V1, V2 and V3, respectively. The smallest displacement was obtained for V2, in which the most material had been added. A different comparison was obtained from displacement projected onto each design’s undeformed CAD surface normals. In V3,




Fig. 6. Redesign strategies and geometric rechecking. (a-d) Original analytical artifact and V1 to V3, viewed from below. The principal envelope was retained in V1, additional warnings that could be addressed geometrically were treated in V2, and material was removed from selected shallow overhangs in V3. Added material, faces exposed by removal and unchanged faces are distinguished by colour. (e,f) Sections through the rear and front halves, with all four design outlines. (g) Block detail with teardrop holes and a pitched channel roof. All 11 reference violations were removed in each redesign. Remaining warnings and dimensional changes are reported in Supplementary Note 6.1.

the 95th percentile of absolute surface-normal deviation was increased from 488 to 537 μm , despite the reduction in maximum displacement. Material savings and shortening of the part were thus accompanied by poorer performance in this dimensional measure. Envelope changes, bore diameters and roundness were also evaluated (Supplementary Notes 6.2 and 6.3). Improvement could not be inferred from maximum displacement alone.

3.6 Process-informed thermo-mechanical redesign

The CMU bracket had no reference-catalog violation, but local stress and plastic-strain concentrations were predicted at the pointed apexes of its web openings and at the web foot, where the web joined the base plate. Three mechanical design alternatives were generated by the agent using this process feedback. The opening apexes and web foot were rounded in S1, the exposed top face of the base was raised by 2 mm in S2, and both changes were combined in S3. In S1, the yielded fraction near the apexes after release was reduced from 5.6% to 1.6%, while maximum whole-part displacement was increased from 0.33 to 0.39 mm. Local strain relief was therefore obtained at the




Fig. 7. Simulated displacement and dimensional deviation of the analytical-artifact redesigns. (a) Maximum displacement after release. The 95th percentile of displacement magnitude is marked in white. (b) The 95th percentile of absolute surface-normal deviation and its maximum divided by ten. (c) Cumulative distributions of absolute surface-normal deviation. Common process settings were used for the original and V1 to V3.

expense of greater global displacement. A geometric rule check alone would not have established this tradeoff (Supplementary Note 6.4).

A separate thermal analysis revealed heat accumulation above the web openings and in the waisted boss, where conduction paths were restricted (See Fig. 8). Accumulation began near 20 mm build height across the absorptivity sensitivity calculations, although absolute temperatures were changed. Four thermal alternatives were generated by the agent and evaluated by simulation. The boss waist was filled to form a straight cylinder in T1. That change was combined with steeper opening roofs in T3 and raised opening floors in T4. Two openings were closed in T2, which was retained only as a reference because features were removed. During reanalysis of an intermediate T3 design, a thermal concentration remained above another opening, prompting a further roof revision. In final T3, the predicted layer-scale peak was reduced from 537 to 377 °C with all four openings retained and 11.1% volume added. Released displacement was increased from 0.33 to 0.37 mm. The same thermal peak was obtained in T4, but a larger released displacement of 0.40 mm was predicted. Thermal improvement was therefore assessed together with material addition, opening geometry and mechanical response (Supplementary Note 7).

3.7 Adaptive support construction through agentic review

Support structures are unavoidable in some cases when geometric redesign is not permitted and the part must be manufactured as designed [4, 25]. The workflow was therefore extended with support design as an alternative to geometric redesign. A clamp model obtained from GrabCAD [31] was used to evaluate this extension, with the original part geometry retained in the supported alternatives. The clamp comprised a split clamping eye and a second boss joined by a web. Its STEP model had a $134 \times 80 \times 50$ mm envelope.

In the initial +Z orientation, five unsupported reaches of 12.45 to 80 mm were identified against the 1 mm horizontal-overhang limit [4]. The addition of triangular ribs beneath the overhangs was evaluated as a geometric remedy. Each rib would have become part of the clamp, with its lower end attached to existing material. However, the adjacent walls ended 11.92 to 15.00 mm above the build plate, leaving no material at the proposed lower attachment locations. The ribs could therefore not be added. No geometric modification was applied in this orientation. A separate coverage limitation was identified during author review: only 34.1% of the downward-facing area was included in the feature-rule checks. Surface checks were therefore added so that downward-facing regions outside the recognized feature classes could also be assessed. Six axis-aligned build directions were screened with these checks, and surfaces still requiring support were identified (Supplementary Note 8.1).

Support geometry was generated for the identified surfaces while the original clamp geometry was retained. Block supports were generated from a grid of 0.2 mm walls and contact teeth, with an initial grid spacing of 1.625 mm [24, 25]. In the initial +Z design, 41,615 mm³ of support was




Fig. 8. Thermal analysis and bracket redesign. (a) Peak-temperature field of the original bracket. (b) Critical regions defined relative to the largest peak, with the two overheating clusters named beside the part. (c-f) Peak-temperature fields for T1, the straight boss, T2, the straight boss with openings 3 and 4 closed, T3, the straight boss with steeper opening roofs, and T4, the straight boss with raised opening floors. T2 removes two features and was retained only as a reference. A common color scale is used for a and c-f. (g-i) Sections through the web of T2, T3 and T4, with the material each design adds to the openings shaded. (j) Section area below the hot cluster, whose height range is shaded. (k) Maximum end-of-scan temperature against build height.

generated before the support catalog existed, and no support criterion was evaluated (Supplementary Fig. S21). However, material was placed beneath steep regions of curved faces that did not require support under the angle criterion. Across the curved targets, 53.5% to 70.6% of each projected footprint lay beneath these regions. The excess material was retained in this initial output.

The coverage limitation and excess support prompted the addition of an Orientation Reviewer, Redesign Reviewer, Support Reviewer and Simulation Reviewer for the respective stage outputs (Fig. 2). The workflow was restarted on the clamp in +Y, -Y, +Z and -X. Only support-catalog review returned work during execution. The other three catalogs were applied after each run for reporting. Grid spacing was reduced in response to measured unsupported spans, and disconnected support pieces were removed. Despite these corrections, support still extended above the curved surface it was intended to carry. The reviewer interfaces were subsequently extended with geometric queries, rendered views and a field for concerns outside the catalogs, so that additional conditions could be investigated (Supplementary Notes 5.3 and 8.2).

The saved designs were then examined by live language-model reviewers. Point classification and rays along the build direction were requested by the Support Reviewer to determine whether part material was present above points inside the support. In +Z, 1,596 mm³, or 3.05% of the support volume, was estimated to have no part surface above it. A sample of 600,000 points was drawn uniformly from the support bounding box and classified against the support solid. Total support volume was reproduced within 0.4% using the same sample. Support wrapping over the crown of the clamping cylinder was identified by this volumetric test, although the projected-footprint criterion had been satisfied.

The defect was traced to the support-envelope code. Selected surface triangles were projected onto the plate and extruded to the highest vertex of their parent face, discarding their local height. On the outer cylindrical target in +Z, the selected triangles reached only 4.7 mm above the plate, whereas the parent face reached 50.0 mm. Support was consequently generated above the intended contact surface. The code was diagnosed and revised by the agent without author guidance during those steps. Each selected triangle was instead swept from its actual position to the plate. In a local comparison, support volume for the cylindrical target was reduced from 24,854 to 1,314 mm³. The flat-target volume was retained at 1,617 mm³, providing a regression control for a surface on which the original height calculation was valid (Fig. 9a and Supplementary Note 8.2).

After local verification, the complete workflow was restarted in the same four build directions with the corrected operator and no human intervention during execution. Support-catalog review was applied during the runs. The other catalogs were applied after run completion, followed by live review of the saved support designs. In +Z, total support volume was reduced from 52,191 to 36,719 mm³, or 29.6%, and the fraction with no part above it was reduced from 3.05% to 0.02% (Fig. 9b,c). The +Z design was subsequently rejected because its exported flat unsupported span of 1.20 mm exceeded the 1 mm limit. The 52,191 mm³ baseline was the later catalog-based design, rather than the initial 41,615 mm³ output. Across the four directions, returns to support construction were reduced from ten to five (Fig. 9d). Reuse of the correction was demonstrated on one clamp in four orientations (Supplementary Note 8.3).

The four supported clamp designs were evaluated by process simulation to estimate distortion after release from the build plate. A mechanical mesh with 1 mm voxels was used, while the support walls were 0.2 mm thick. Individual support walls could not be resolved by this mesh, so the predicted displacements were treated as preliminary estimates. Maximum displacements of 0.830, 1.015, 0.779 and 0.970 mm were predicted for +Y, −Y, +Z and −X, respectively (Fig. 9i-l).

In +Z, the unsupported span increased from 1.00 mm on the in-loop solid to 1.20 mm on the exported geometry, although both checks used 0.15 mm sampling. The operation responsible for this discrepancy remained unresolved. The exported measurement exceeded the 1 mm limit, and the design was rejected. This discrepancy demonstrates the need to recheck the saved geometry before assigning an acceptance decision. In −X, a span of 1.12 mm was measured only on curved surfaces. That alternative was retained with a warning because the flat-overhang limit had been extrapolated to curved geometry. Three alternatives were accepted with warnings and +Z was rejected in the support-stage review. Surface checks still identified powder-contact violations below 25° in all four designs. In −X, an orientation-comparison violation also remained because +Y and −Y were dominating alternatives under ORI-C1.

For the AgenticDfAM evaluation, a material-use warning was assigned when support volume exceeded 25% of part volume. The −X alternative required support equal to 39.8% of part volume and was retained with this warning. The support-tooth dimensions were taken from experiments on IN625 [25] and used in the clamp simulations. Support removability was not evaluated. An applicability warning was also retained because the contact-tooth dimensions had been adopted from a different material, without experimental validation on IN718 (Supplementary Note 8.3).




Fig. 9. Verification of the corrected support operator, supported clamp designs and predicted displacement after plate release. (a) Test of the corrected operator on two single faces before it was reused in the complete runs. On a flat underside, where the original operator was already correct, the support volume was unchanged at 1,617 mm^3 . On a curved underside, where the defect occurred, it was reduced from 24,854 to 1,314 mm^3 . (b,c) The +Z design before and after the envelope correction, viewed from one camera. Support triangles whose centroid lies above the highest part surface at the same plan position are shown in crimson. Before the correction, a slab of support rested on the crown of the horizontal tube, and 3.05% of the support volume had no part above it. After the correction, this fraction was 0.02%, and in -X it was reduced from 0.31% to zero. (d) Number of times the support reviewer rejected a support design and sent it back to support construction for another attempt, in each of the four complete reruns. The total was reduced from ten to five. (e-h) Final part and support geometries, viewed from below in the +Y, -Y, +Z and -X build frames, with the support volume before and after the correction. (i-l) Displacement magnitude after plate release on a common color scale, with the support from which each part was released ghosted in orange. Maximum values of 0.830, 1.015, 0.779 and 0.970 mm were predicted. The pre-correction -X simulation exceeded its eight-hour budget. The corrected reruns used the final support catalog. The lowest maximum displacement was predicted for +Z, but on review of the exported +Z geometry a flat unsupported span of 1.20 mm was measured against the applicable 1 mm limit, and the design was rejected. The fields were treated as exploratory predictions because 0.2 mm support walls were not resolved by the 1 mm mechanical voxels. Pitches, measured spans and recorded support-stage decisions are given in Supplementary Tables S21 and S22.

4 Discussion

AgenticDfAM connects accurate geometric measurement with reasoning about manufacturing guidance, design alternatives and process response. The resulting engineering capability is the revision of a design together with the investigations and construction procedures used to assess and construct it. In the architecture of Fig. 1, analyses and redesign actions are selected by agents, while geometric quantities, rule comparisons and CAD operations are supplied by tools. Reasoning is thereby connected to executable changes, and engineering judgments remain tied to numerical evidence. Adaptation can be directed at the geometry, the interpretation of a criterion or the procedure used to construct a candidate design.

Accurate measurements provide the numerical basis for reasoning, while feature context and source applicability determine how those measurements should be judged. Agreement with reference dimensions coexisted with limitations in bore recognition and wall-support classification. Different compliance judgments were also obtained from unchanged geometry when recommendations and experimental limits were interpreted differently. Assessment therefore had to connect the measured feature to the conditions supporting the guideline and the requirements of the intended product. Retaining the measurement and its applicability rationale together allows an engineering judgment to be examined or revised without altering the numerical evidence.

The redesign cases demonstrate why reasoning must extend beyond violation removal. Alternatives satisfying the same geometric criteria incurred different material and dimensional consequences. Process feedback also enabled changes beyond those triggered by the geometric catalog: the bracket satisfied the reference rules, yet predicted stress and thermal concentrations were used by the agent to develop further alternatives. Reanalysis prompted an additional roof revision when a thermal concentration remained. Local yielding and heat accumulation were reduced in the predictions, while displacement after plate release increased. Iterative redesign could thus be directed by process evidence even when geometric checks had been passed. Selection must still be tied to product requirements because geometric compliance and feature retention do not establish functional equivalence.

Procedural adaptation was demonstrated in support construction, where an unexpected geometric outcome was identified despite satisfaction of the available support checks. Additional volumetric queries requested by the Support Reviewer exposed support above its intended contact surface, and the construction code was diagnosed and corrected by the agent. Verification on a curved target and a flat control established the intended geometric effect before reuse. The projected-footprint criterion was passed before and after the correction. The investigation therefore changed how a design was constructed while retaining the numerical evidence used to assess it.

5 Conclusions

AgenticDfAM advances DfAM by making the workflow itself adaptable. Design guidance, geometric recognition and measurements, as well as process-informed feedback are connected through reasoning agents that direct executable CAD, analysis, support generation, and redesign tools. Assessment can be extended when prescribed checks are insufficient, additional evidence can be requested, and both designs and construction procedures can be revised. Geometric measurements and numerical checks are performed by the developed Python tools, while the results are interpreted by the agents to guide assessment, redesign and further investigation.

This capability was demonstrated across geometric assessment, redesign and process simulation. Agreement with the NIST reference geometric quantities at the reporting resolution established the measurement basis. Three CAD redesign strategies removed the analytical artifact’s reference violations while exposing different material and distortion consequences. For the CMU bracket, process evidence was interpreted by the agent to generate and refine alternatives, reducing the predicted layer-scale peak temperature by 160 °C while material addition and released displacement were assessed together. Adaptation extended to the construction code when additional geometric

queries enabled a support-generation procedure to be diagnosed, corrected and verified before reuse across four build orientations.

The broader contribution is a framework through which manufacturing knowledge can be translated into engineering actions and returned evidence can reshape those actions. New concerns can be investigated beyond the initial rule catalog, and the procedures used to produce a design can themselves become subjects of engineering revision. AgenticDfAM thus provides a foundation for expanding DfAM automation through reasoning grounded in accurate measurement, with the evidence behind each design and procedural change retained for inspection and verification.

References

- [1] M.K. Thompson, G. Moroni, T. Vaneker, G. Fadel, R.I. Campbell, I. Gibson, A. Bernard, J. Schulz, P. Graf, B. Ahuja, F. Martina, Design for Additive Manufacturing: Trends, opportunities, considerations, and constraints, *CIRP Annals* 65 (2016) 737–760. <https://doi.org/10.1016/j.cirp.2016.05.004>.
- [2] G.A.O. Adam, D. Zimmer, On design for additive manufacturing: evaluating geometrical limitations, *Rapid Prototyp. J.* 21 (2015) 662–670. <https://doi.org/10.1108/rpj-06-2013-0060>.
- [3] J. Kranz, D. Herzog, C. Emmelmann, Design guidelines for laser additive manufacturing of lightweight structures in TiAl6V4, *J. Laser Appl.* 27 (2015) S14001. <https://doi.org/10.2351/1.4885235>.
- [4] D. Herzog, K. Asami, C. Scholl, C. Ohle, C. Emmelmann, A. Sharma, N. Markovic, A. Harris, Design guidelines for laser powder bed fusion in Inconel 718, *J. Laser Appl.* 34 (2022) 012015. <https://doi.org/10.2351/7.0000508>.
- [5] H.D. Budinoff, S. McMains, Will it print: a manufacturability toolbox for 3D printing, *International Journal on Interactive Design and Manufacturing (IJIDeM)* 15 (2021) 613–630. <https://doi.org/10.1007/s12008-021-00786-w>.
- [6] Y. Shi, Y. Zhang, S. Baek, W. De Backer, R. Harik, Manufacturability analysis for additive manufacturing using a novel feature recognition technique, *Comput. Aided. Des. Appl.* 15 (2018) 941–952. <https://doi.org/10.1080/16864360.2018.1462574>.
- [7] P. Stavropoulos, K. Tzimanis, T. Souflas, H. Bikas, Knowledge-based manufacturability assessment for optimization of additive manufacturing processes based on automated feature recognition from CAD models, *The International Journal of Advanced Manufacturing Technology* 122 (2022) 993–1007. <https://doi.org/10.1007/s00170-022-09948-w>.
- [8] M. Mayerhofer, W. Lepuschitz, T. Hoebert, M. Merdan, M. Schwentenwein, T.I. Strasser, Knowledge-Driven Manufacturability Analysis for Additive Manufacturing, *IEEE Open Journal of the Industrial Electronics Society* 2 (2021) 207–223. <https://doi.org/10.1109/ojies.2021.3061610>.
- [9] K. Hu, S. Jin, C.C.L. Wang, Support slimming for single material based additive manufacturing, *Computer-Aided Design* 65 (2015) 1–10. <https://doi.org/10.1016/j.cad.2015.03.001>.
- [10] C. Ellsel, R. Stark, A knowledge-driven, integrated design support tool for additive manufacturing, *Proceedings of the Design Society* 4 (2024) 1747–1756. <https://doi.org/10.1017/pds.2024.177>.
- [11] S. Jalui, J. O’Connor, Y. Xuan, G. Manogharan, Developing novel restrictive design for additive manufacturing (DfAM) constraints for NURBS-based adjoint shape optimization for metal AM, *Graph. Models* 142 (2025) 101304. <https://doi.org/10.1016/j.gmod.2025.101304>.
- [12] Autodesk, To Display Simulation Results into Netfabb, (2020). <https://help.autodesk.com/cloudhelp/2020/ENU/NETF-Workflows/files/GUID-9A971083-A2BE-401F-82E1-2513788C8FBA.htm>.
- [13] M. Langelaar, Combined optimization of part topology, support structure layout and build

- orientation for additive manufacturing, *Structural and Multidisciplinary Optimization* 57 (2018) 1985–2004. <https://doi.org/10.1007/s00158-017-1877-z>.
- [14] M.R. Correa, C.-J. Thore, R.F. Ausas, S. Jakobsson, G.A. Haveroth, J.A. Cuminato, A transient thermoelastic mathematical model for topology optimization of support structures in additive manufacturing, *Structural and Multidisciplinary Optimization* 67 (2024) 35. <https://doi.org/10.1007/s00158-024-03757-3>.
 - [15] P.K. Jayaraman, A. Sanghi, J.G. Lambourne, K.D.D. Willis, T. Davies, H. Shayani, N. Morris, UV-Net: Learning From Boundary Representations, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021: pp. 11703–11712. https://openaccess.thecvf.com/content/CVPR2021/html/Jayaraman_UV-Net_Learning_From_Boundary_Representations_CVPR_2021_paper.html.
 - [16] J.G. Lambourne, K.D.D. Willis, P.K. Jayaraman, A. Sanghi, P. Meltzer, H. Shayani, BRepNet: A Topological Message Passing System for Solid Models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021: pp. 12773–12782. https://openaccess.thecvf.com/content/CVPR2021/html/Lambourne_BRepNet_A_Topological_Message_Passing_System_for_Solid_Models_CVPR_2021_paper.html.
 - [17] H. Wu, R. Lei, Y. Peng, L. Gao, AAGNet: A graph neural network towards multi-task machining feature recognition, *Robot. Comput. Integr. Manuf.* 86 (2024) 102661. <https://doi.org/10.1016/j.rcim.2023.102661>.
 - [18] Y. Zhang, Y.F. Zhao, Hybrid sparse convolutional neural networks for predicting manufacturability of visual defects of laser powder bed fusion processes, *J. Manuf. Syst.* 62 (2022) 835–845. <https://doi.org/10.1016/j.jmsy.2021.07.002>.
 - [19] M.T. Khan, W. Feng, L. Chen, Y.H. Ng, N.Y.J. Tan, S.K. Moon, Automatic Feature Recognition and Dimensional Attributes Extraction From CAD Models for Hybrid Additive-Subtractive Manufacturing, in: *Volume 2A: 44th Computers and Information in Engineering Conference (CIE)*, 2024: p. V02AT02A035. <https://doi.org/10.1115/detc2024-143107>.
 - [20] A.C. Doris, D. Grandi, R. Tomich, M.F. Alam, M. Ataei, H. Cheong, F. Ahmed, DesignQA: A Multimodal Benchmark for Evaluating Large Language Models’ Understanding of Engineering Documentation, *J. Comput. Inf. Sci. Eng.* 25 (2025) 21009. <https://doi.org/10.1115/1.4067333>.
 - [21] M.T. Khan, L. Chen, Y.H. Ng, W. Feng, N.Y.J. Tan, S.K. Moon, Leveraging Vision-Language Models for Manufacturing Feature Recognition in Computer-Aided Designs, *J. Comput. Inf. Sci. Eng.* 25 (2025) 104501. <https://doi.org/10.1115/1.4069266>.
 - [22] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, Y. Cao, ReAct: Synergizing reasoning and acting in language models, in: *International Conference on Learning Representations (ICLR)*, 2023. <https://doi.org/10.48550/arXiv.2210.03629>.
 - [23] Open Cascade SAS, Open CASCADE Technology 7.9.3 with the OCP Python bindings of the CadQuery project, (n.d.). <https://dev.opencascade.org>.
 - [24] A. Dimopoulos, M. Salimi, T.-H. Gan, P. Chatzidakos, Support Structures Optimisation for High-Quality Metal Additive Manufacturing with Laser Powder Bed Fusion: A Numerical Simulation Study, *Materials* 16 (2023) 7164. <https://doi.org/10.3390/ma16227164>.
 - [25] Z. Li, G. Fei, D. Wu, X. Chen, Y. Yu, Z. Liu, J. Zhang, J. Zhou, Optimization of Parameters of Block-Shaped Support Tooth Structure Using Orthogonal Experimental Design in Laser Powder Bed Fusion, *Materials* 19 (2026) 1480. <https://doi.org/10.3390/ma19081480>.
 - [26] J. Trapp, A.M. Rubenchik, G. Guss, M.J. Matthews, In situ absorptivity measurements of metallic powders during laser powder-bed fusion additive manufacturing, *Appl. Mater. Today* 9 (2017) 341–349. <https://doi.org/10.1016/j.apmt.2017.08.006>.
 - [27] R. Ranjan, C. Ayas, M. Langelaar, F. van Keulen, Fast detection of heat accumulation in powder bed fusion using computationally efficient thermal models, *Materials* 13 (2020) 4576.

<https://doi.org/10.3390/MA13204576>.

- [28] R. Ranjan, C. Ayas, M. Langelaar, F. van Keulen, Controlling local overheating in topology optimization for additive manufacturing, *Structural and Multidisciplinary Optimization* 65 (2022) 162. <https://doi.org/10.1007/S00158-022-03258-1>.
- [29] D. Guirguis, C.S. Tucker, J.L. Beuth, Machine learning for real-time detection of local heat accumulation in metal additive manufacturing, *Mater. Des.* 241 (2024). <https://doi.org/10.1016/j.matdes.2024.112933>.
- [30] S. Moylan, J. Slotwinski, A. Cooke, K. Jurrens, M.A. Donmez, An Additive Manufacturing Test Artifact, *J. Res. Natl. Inst. Stand. Technol.* 119 (2014) 429–459. <https://doi.org/10.6028/jres.119.017>.
- [31] À. Maynegre, Split shaft clamp bracket with offset pivot, (2026). <https://grabcad.com/library/split-shaft-clamp-bracket-with-offset-pivot-1> (accessed April 30, 2026).
- [32] P.S. Foundation, The Python Language Reference, (2024). <https://docs.python.org/3.13/reference/>.
- [33] C. contributors, OCP: Python bindings for Open CASCADE Technology, (2026). <https://github.com/CadQuery/OCP>.
- [34] W. Schroeder, K. Martin, B. Lorensen, *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, 4th ed., Kitware, Clifton Park, NY, 2006.
- [35] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N.J. Smith, R. Kern, M. Picus, S. Hoyer, M.H. van Kerkwijk, M. Brett, A. Haldane, J.F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, T.E. Oliphant, Array programming with NumPy, *Nature* 585 (2020) 357–362. <https://doi.org/10.1038/s41586-020-2649-2>.
- [36] J.D. Hunter, Matplotlib: A 2D Graphics Environment, *Comput. Sci. Eng.* 9 (2007) 90–95. <https://doi.org/10.1109/MCSE.2007.55>.
- [37] S. Colvin, E. Jolibois, H. Ramezani, A. Garcia Badaracco, T. Dorsey, D. Montague, S. Matveenko, M. Trylesinski, S. Runkle, D. Hewitt, A. Hall, V. Plot, Pydantic Validation, (2026). <https://github.com/pydantic/pydantic>.
- [38] American Welding Society, Specification for Fabrication of Metal Components using Additive Manufacturing, (2019). https://pubs.aws.org/Download_PDFS/AWS_D20_1_D20_1M_2019-PV.pdf (accessed October 17, 2026).
- [39] U.S. Department of Defense, Technical Data Packages, (2025). https://quicksearch.dla.mil/qsDocDetails.aspx?ident_number=276980 (accessed May 17, 2026).
- [40] National Aeronautics and Space Administration, Additive Manufacturing Requirements for Spaceflight Systems, (2021). <https://standards.nasa.gov/standard/NASA/NASA-STD-6030> (accessed May 17, 2026).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [AgenticDfAMSupplementaryInformationPreprint.pdf](#)