Screening STEP-to-Web Conversion Errors

An Engineering Technical Report on Local Sampling, Source Completeness and Independent Comparisons

Shiya Yuan
BYBOLT

22 October 2026 | v0.11 technical report
Not peer reviewed

DOI: https://doi.org/10.5281/zenodo.23554447

Abstract

A coarse sampled-distance gate can accept inaccurate CAD-derived web meshes. We extend an earlier conversion case study with local geometric probes, a query-count-matched random control and an independently implemented surface mesher. The new screen probes triangle centroids, indexed-edge midpoints, each CAD face and CAD edges, alongside a small random sample. It detects all three previously exposed gate misses and preserves three exposed passing controls. An initial twenty-file test exposes silent first-root-only import of a three-root assembly. After correcting every used reference/export path, that cohort is retained as amendment evidence and the policy is frozen before importing twelve further file-disjoint community CAD files. Eleven files are measurable and one is empty. All three sampling methods detect the single observed violating baseline; none misses an observed violation in this final cohort. Gmsh is attempted as a fixed alternative mesher, with source failures and resource caps retained. Trimesh, explicit facewise CAD-distance minima and a separate Gmsh import/binding cross-check selected witnesses. These strengthen implementation checks but share the OpenCASCADE geometry family. The result evaluates sampled detection and engineering reproducibility; it does not establish continuous error bounds, a new general meshing algorithm, or independent-kernel truth. Detecting or withholding an output is not geometric repair.

1 Purpose, audience and prior evidence

This technical report is intended for engineers maintaining CAD-to-web conversion and verification workflows. It documents an implementation failure analysis, a reproducible screening procedure and its measured limitations. Revisions v0.9–v0.11 refine presentation, literature positioning and publication declarations; the numerical evidence and frozen experimental policies remain those of v0.8/v08b. This preprint reports completed, bounded engineering observations and has not been peer reviewed.

Successful web loading and correct overall dimensions do not establish local correspondence to a source CAD boundary. This study asks whether local geometric coverage reduces missed surface deviations, and whether any gain survives a control using the same number of distance queries. A separate comparison asks how a different mesh generator behaves under fixed settings. We distinguish detection, preservation, abstention and geometric improvement throughout.

Earlier work in this project repaired two product meshes sharing a threaded rod by partitioning long B-spline faces. Later synthetic and external tests exposed pose dependence, export rejection and missed deviations. Skipping unnecessary sewing resolved a real regression: in the previous twenty-file cohort v01–v20, all seventeen distance-passing baselines were retained byte-for-byte, recovering five old-repair rejections. However, the 512-point gate accepted three outputs that failed different-seed evaluation. No baseline outlier in that cohort was corrected. That earlier result remains evidence of preservation, not successful general repair. The newly discovered multi-root issue also limits earlier statements about complete assembly import; those historical measurements concern the shapes actually loaded, unless source completeness is separately established.

The three misses are v11/v12, T-nuts with source labels M3/M4, and v15, a split washer labelled M20. Their old-gate maxima were 0.021619, 0.043503 and 0.021490 mm against the study's 0.05 mm budget.

They are now exposed development cases, not held-out evidence. Prior topology/normal diagnostics were zero on these three, so reinstating a hard diagnostic veto would not solve them. Six other outputs in the earlier cohort retained diagnostic flags; sampled surface accuracy never implied topology certification.

2 Prior practice and the report's scope

The literature check is targeted, not systematic. Bibliographic identity, access level, relevant sections and the claims supported by each reference are recorded in `release/v09/literature-audit.json`. No third-party figures, tables or article PDFs are reproduced in the release candidate.

Parameterized fastener construction is established: Zhang describes a SolidWorks/Visual Basic hex-bolt workflow in the publisher abstract [2]; Bozherikov describes design-table components and a constrained bolted assembly [3]; cq_warehouse supplies parameterized fasteners in software [4]. These sources establish context, not a measured baseline for this report. The CAD-quality survey distinguishes geometric quality from broader modeling and semantic concerns [5]; a sampled surface pass here does not establish design intent or fitness for manufacture.

CAx-IF already specifies volume, area and centroid checks (Section 4), face-associated sampling through COPS (Section 6), and assembly child-count and placement validation (Section 7) [6]. Source completeness and local sampling are therefore established validation concerns. Our imported-root/solid counts are not the STEP product-structure child counts specified there, and our detector does not implement COPS exchange encoding or claim CAx-IF conformance. Its practical contribution is the retained failure, correction and controlled evidence in this particular workflow.

Sap and Szabo's accessible 2020 preprint describes property-conditioned interchangeability and round-robin checks across CAD systems [8]. Their 2025 journal record describes representation-independent local/global comparisons and user-selected tolerances [7]. The latter was checked through publisher-indexed abstract, introduction and section snippets; full text was not obtained. The earlier preprint is related work, not a verified substitute for the journal article. Neither framework was reproduced here, so no comparative performance or general novelty claim follows.

Li et al. address intersection-free tessellation with critical regions and constrained triangulation [9]. Zhou et al.'s preprint prioritizes topology-preserving mesh construction [10]; its Section 5.1 explicitly evaluates the first part of multi-part models. That declared evaluation scope differs from the silent source truncation found here and is not evidence of a defect in their work. Both papers concern mesh generation and stronger geometric/topological objectives; this report evaluates sampled screening and import behavior, without reimplementing either method or borrowing their guarantees.

CadQuery and glTF document exchange and coordinate conventions [11, 12]. Gmsh supplies a separate mesher and import binding [16, 15]; Trimesh supplies a separate mesh parser and triangle-distance implementation [13]. These implementation comparisons do not establish independence from the shared OpenCASCADE geometry family. Online software documentation is version-sensitive; executed versions and frozen source hashes, rather than current documentation alone, define the experiment.

3 Independent comparison exposed incomplete source import

The first twenty-file detector cohort contained two observed baseline deviations among nineteen single-root sources: w09/w10, split washers labelled M16/M30. The old sampler missed w09; both local and matched sampling detected both, and the separate evaluator confirmed them. The twentieth source, bearing w20, is not a valid complete-source accuracy comparison under that implementation.

Gmsh appeared to deviate by about 2.5 mm on w20, but an independent import/distance check disagreed with that conclusion. A post-outcome audit found three STEP root objects. The earlier CadQuery code used `importStep(...).val()`, retaining only the first. It therefore exported and evaluated one solid with four faces and volume 50.265482 mm^3. Gmsh imported three solids with fourteen faces and summed volume 276.987939 mm^3. The purported mesher regression was a reference-completeness error. It is retained as a failed comparison, not attributed to Gmsh geometry quality.

The amended loader collects every imported shape using `vals()` and constructs a compound without changing placements. Baseline export, reference sampling, screening and facewise witness checks all use

that loader. On exposed w20 it recovers three solids/fourteen faces and agrees with Gmsh volume within 10^{-8} mm^3. The retained partial mesh and fresh complete export are tested separately. Audits of the first cohort, all fifty-four earlier external files and all thirty-three source-library STEP files find w20 is the only multi-root case among those inspected files. The library files are read only; their resource bytes remain unchanged.

This correction changes what geometry is evaluated, even though the sampling rules are unchanged. We therefore do not relabel the corrected first cohort as untouched validation. Twelve further files are selected after the amendment, with all seventy-four earlier external file identities excluded. The shared source-completeness defect demonstrates why agreement between algorithms over the same incomplete reference is insufficient.

4 Methods

4.1 Distance and acceptance semantics

The reference is the original STEP boundary as imported by the recorded CAD kernel, not an approved manufactured part. Mesh-to-CAD distance is measured to the compound of CAD shells, including for points inside a solid. CAD-to-mesh distance is measured to triangles, not just vertices. The observed error is the largest queried distance in either direction. The 0.05 mm budget is an illustrative study choice; it is not a dimensional standard, fit requirement, manufacturing tolerance or validated visual threshold.

The original gate uses 512 samples per direction, up to 64 face probes and up to 64 flagged-triangle centroids. The new screen additionally runs the local detector below. An observed over-budget point produces rejection. Missing, capped or failed checks produce an inconclusive result, never acceptance. Only when both checks pass is an unchanged accepted copy written. Diagnostic flags remain attached and require separate review. All outcomes remain finite-sampling statements, without a Hausdorff upper bound or a topology certificate.

4.2 Local geometric coverage

The local mesh-to-CAD probe set contains 512 area-weighted random points, every exported triangle centroid and every unique indexed-edge midpoint. Edge identity uses mesh indices, so distinct indices on geometrically coincident seams can produce duplicate positions; all queries are counted. We do not target stored failure coordinates, part names or manually selected bounding boxes.

The CAD-to-mesh set contains 512 area-weighted reference points projected to the trimmed CAD boundary. In addition, each CAD face contributes centroids from up to 64 evenly indexed triangles in its reference tessellation, individually projected onto that face. Each CAD edge contributes nine normalized-length positions, including endpoints. This prevents large faces from receiving all the attention solely because of area. It still does not cover every continuous point, and the face probes depend on reference triangulation order. An initial development probe using only triangle centroids missed v11; edge midpoints and CAD-face probes are therefore retained, with failed exploratory outputs preserved.

Local limits are 100,000 mesh queries and 20,000 additional CAD probes. Exceeding either limit is inconclusive, not a clean mesh. The full detector can be expensive and is not a real-time browser feature.

4.3 Matched random control and separate evaluation

For each baseline, a global area-weighted random control receives exactly the same number of mesh-to-CAD and CAD-to-mesh queries as the local detector. Its seeds differ. It has no per-face, per-edge or triangle-centroid allocation. The match concerns distance-query counts, not point-generation cost, wall time or spatially paired points. A local resource failure prevents construction of a complete matched count and is disclosed.

A separate evaluator uses 8192 random points per direction and other fixed seeds. Its point sets are created only after the detector decisions are saved. Seeds are 101001/102001 for local CAD/mesh samples, 103001/104001 for matched random and 105001/106001 for evaluation. Reference samples are projected from an existing CAD tessellation, so the evaluator is not wholly independent of the CAD representation.

A verified over-budget witness from any stage is evidence of a violation even if another sampler misses it. We report misses against the union of observed baseline violations and separately identify those seen by the 8192-point evaluator. This union is not exhaustive ground truth, and its sensitivity over an industrial population is unknown. Passing samples cannot disprove a larger valid distance found elsewhere.

4.4 Independent mesh generation and numerical checks

Every final held-out source is also attempted with Gmsh 4.15.2 [15]. The fixed settings are surface algorithm 6, first-order triangles, 48 curvature elements per full circle, 0.02 mm minimum size, and maximum size $\min(1.5, \max(0.25, D/25))$ mm for bounding-box diagonal D. The run uses one thread and seed 108001. Coordinates are converted from mm/Z-up to m/Y-up for GLB. The source STEP is unchanged. This comparator has its own mesh sizes and triangle counts; it is not matched in bytes or triangles and is never silently selected as a repair fallback.

Three checks address implementation error. First, Trimesh independently parses each measured GLB and runs exhaustive point-to-triangle queries. Second, the compound CAD distance is compared with the minimum of explicit distances to every imported face. Each uses up to eight index-spaced and eight worst retained witnesses in the relevant direction. Third, Gmsh separately imports the STEP and evaluates its own OCC point-to-face distances on up to four index-spaced and four worst mesh witnesses. The tolerances for implementation agreement are 10^{-7} mm for the first two checks and 10^{-6} mm for the third.

Gmsh's mesh generator and binding/import route are separate implementations, but its OCC 7.8.1 and CadQuery's OCP 7.8.1.1.post1 belong to the same geometry-kernel family. Agreement is not independent-kernel confirmation or a physical measurement. Witness selection verifies numerical maxima; it is not an unbiased population sample.

5 Protocol and source coverage

The initial development group comprised three old misses and passing controls v01/v06/v13. A first frozen detector was tested on twenty new files w01–w20. Its independent comparison uncovered incomplete STEP-root import on w20, requiring an implementation amendment. That entire first cohort is retained as exposed amendment evidence, not final held-out evidence. The final development group includes v11/v12/v15, passing v06, the retained partial w20 mesh and a fresh complete w20 export. The importer and all source hashes are frozen again before the final twelve sources are imported. Sampling settings, thresholds and Gmsh meshing parameters are unchanged. No development or first-cohort result is pooled with final held-out counts.

Twelve further new file identities x01–x12 are selected from the same pinned FreeCAD community-library revision f5ca01cdcaa2f69b9295a6382719c2d90628780c [1]. The selection and source hashes are retained before geometry import. All twelve Git blobs are distinct and absent from the earlier 74 external source files. Selection uses filenames and intended family coverage: six files labelled as bearings, two screws, three split washers and one retaining ring. No failed file is replaced. Source x06 is an empty Git blob in the pinned repository; exact blob verification distinguishes it from a truncated network transfer. It remains a requested input failure, not a valid CAD model. The transport-only resume retains these exact bytes without altering the frozen detector or replacing its identity.

This is purposeful file-disjoint selection, not a random industrial sample. Families, authors, templates and repository overlap. No unused T-nut or metric-nut STEP files remain in this pinned library, so held-out T-nut family transfer is not established. Source filenames are labels, not verified standards, dimensions, materials or performance. Assets retain CC BY 3.0 with per-file contributor history.

The GLB cap is 16 MiB and 750,000 indexed triangles, checked before loading but after export. This is not a peak-memory guarantee. Baseline export is limited to 90 seconds, Gmsh export to 300 seconds, new numerical stages to 600 seconds, and additional Gmsh witness batches to 120 seconds. The old gate retains its 180-second limit. Two case workers make recorded timings unsuitable for a controlled runtime ranking. Incomplete results remain in the twelve-file denominator.

6 Results

6.1 Exposed regressions

The local detector and matched random control both detect the three known surface misses. The three initial passing controls remain below budget. The amended development table additionally separates a deliberately retained partial assembly from a fresh complete export. Actual screen-entry-point runs reject the known bad outputs; a passing control is copied byte-for-byte. This verifies screening and import correction, not geometric repair of the three original surface errors.

Case	Old gate	Local	Matched	Evaluation	Gmsh maximum
v11	0.021619	0.062413	0.062183	0.061328	0.062384
v12	0.043503	0.087314	0.087111	0.086591	0.087313
v15	0.021490	0.059406	0.057178	0.058335	0.027629
v06	0.001158	0.001173	0.001163	0.001158	0.004314
w20-partial	2.500777	2.500777	2.500777	2.500777	0.010847
w20-full	0.001554	0.001554	0.001554	0.001554	0.010847

All distances are mm. In this amended table, "Old gate" means the old sampling rule using the corrected complete source; it is not the legacy incomplete-reference result on w20. The two w20 rows are different exports of one source, not independent CAD models. Gmsh maximum combines its local and separate-evaluation observations. Gmsh improves the exposed washer below budget but leaves both T-nuts over budget. The development evidence thus supports detecting these regressions, while failing to establish superiority of local probes over equally numerous random samples.

6.2 Twelve unexposed files

The observed baseline violations are x11; those seen by the separate dense evaluator are x11. Counts below distinguish incomplete checks from missed violations.

Detector	Measured	Detected	Observed misses	Incomplete
Old gate	11	1	0	1
Local probes	11	1	0	1
Matched random	11	1	0	1

The combined screen rejects one file, leaves the empty input inconclusive and passes the other ten files. Rejection prevents a sampled pass; it does not reduce the underlying mesh error. Local query totals range from 6942 to 17944, with median 12015 per baseline. These counts are much larger than the earlier small gate and must be considered when interpreting improvement.

ID	Old gate	Local	Matched	Evaluation	Gmsh maximum
x01	0.003419	0.003419	0.003419	0.003419	0.023250
x02	0.002020	0.002020	0.002020	0.002020	0.015423
x03	0.004973	0.004973	0.004973	0.004973	0.030038
x04	0.004041	0.004041	0.004041	0.004041	0.026819
x05	0.001785	0.001787	0.001787	0.001787	0.011971
x06	—	—	—	—	—*
x07	0.001683	0.001864	0.001688	0.001755	0.008633
x08	0.003720	0.003729	0.003729	0.003729	0.011309*
x09	0.016188	0.022011	0.020538	0.020675	0.010844
x10	0.023594	0.031289	0.030865	0.030676	0.022741
x11	0.064140	0.068194	0.066637	0.067650	0.029065
x12	0.004365	0.004392	0.004391	0.004392	0.021763

Values are mm. A dash denotes a missing measurement. An asterisk indicates that the Gmsh local/evaluation pair is incomplete; its available maximum is not a complete pass. Gmsh violations occur in none; incomplete comparator cases are x06, x08. Baseline violating cases with a complete below-budget Gmsh result are x11. New observed Gmsh violations on baselines without an observed violation are none. Source x06 is empty; x08 exceeds the local probe cap on its Gmsh mesh despite a below-budget separate evaluation. These are fixed-setting observations, not a fair resource-matched mesher ranking.

6.3 Source-truncation controls and coverage gap

A separate ablation was fixed before final-cohort geometry import: every multi-root source would be subjected to a deliberately first-root-only export. In fact, all eleven valid final files import as a single root and a single solid, despite five bearing labels. The empty x06 source is an input failure. Thus zero final sources are eligible, no injected omission trials run, and this planned control supplies no held-out evidence for multi-root generalization. A source filename does not establish assembly complexity.

The exposed w20 regression remains the real multi-root example. As a separate post-outcome implementation check, three synthetic multi-root STEP fixtures contain a box, cylinder and sphere: original placement, rigid rotation/translation, and permuted root order. Analytic total volume and independently specified component centers are compared with both import paths. All three fixtures retain three roots/solids and agree within the predeclared 10^{-7} tolerances; the legacy first-root access retains one solid. These simple controls verify import/order/placement behavior, not transfer to unseen real assemblies or curved-surface meshing fidelity. A NumPy-boolean JSON serialization failure is retained, followed by a serialization-only repair and rerun with unchanged geometry and checks.

6.4 Independent checks and integrity

The held-out numerical checks include 864 explicit facewise CAD witnesses, 862 Trimesh triangle-distance witnesses and 432 additional Gmsh import/distance witnesses. Their largest discrepancies are respectively 0, 6.33e-14 and 2.21e-13 mm. All completed witness checks satisfy their declared agreement tolerances. The underlying arrays, errors and statuses are retained, rather than replaced by aggregate success labels.

All 132 protected source-library and storefront STEP/GLB asset hashes remain unchanged. The earlier conditional-sewing helper is unchanged in this revision. The new screen, experiment scripts and paper are local research changes. During the numerical study, no deployment, browser performance test or public submission was performed. Frozen-source/input staging is checked in a new directory; a second complete numerical replay of v0.8 or another operator's reproduction is not claimed.

7 Interpretation and limitations

The exposed misses are resolved as detection failures: the new screen refuses all three. On the new files, the old gate misses 0 observed violations, local probes miss 0 and matched random misses 0. These counts, along with incomplete checks, determine the supported scope. The local and matched controls have the same detection outcomes here, so this study supplies no evidence that geometric stratification itself improves classification at that query budget. The final twelve-file cohort also contains no small-gate miss, so it does not independently demonstrate a detection advantage over the old sampler. It may still locate useful witnesses, but greater maxima alone do not demonstrate better operational decisions.

No finite collection of centroids, midpoints, edge positions or random points certifies continuous surface error. The reported observed-violation union can also miss defects. Source-family dependence, purposeful selection and the absence of new T-nut examples limit generalization; no binomial confidence interval treating all files as independent industrial parts is justified. Diagnostic flags and topology remain separate concerns. Manufacturing standards, thread fit, PMI preservation, human visual quality and mechanical performance are outside the claims.

The independent Gmsh path has already exposed a source-import defect that same-reference checks missed. The fixed Gmsh comparison also improves the evidence beyond testing one exporter repeatedly. On the eleven valid final sources, both paths report the same solid counts and their volumes differ by at most 7.74×10^{-12} mm^3; this basic agreement is useful but insufficient to prove boundary equivalence. Its resource caps and remaining violations also prevent a universal replacement claim. Both CAD paths retain common kernel ancestry, and explicit facewise checks can share kernel-level errors. Independent-kernel or approved-reference comparisons and another person's reproduction remain necessary for stronger claims.

This report contributes a reproducible engineering failure analysis: a shared incomplete reference can hide a conversion defect, while an independent import path can expose it. It also records a negative controlled result: local probes do not outperform query-count-matched random sampling on the measured classification outcomes. This is useful operational evidence without establishing a new general validation

principle. A stronger research-paper submission would require fuller comparison with the closest literature, deployment-independent reproduction by another operator, fresh multi-source and multi-root cases, and evidence appropriate to any claimed algorithmic improvement. Those experiments are not prerequisites for describing the present bounded observations honestly in a technical report.

8 Practical use and evidence map

For a local engineering investigation, first retain the source bytes and expected content scope. Record imported roots, solids and placements, and use an independently implemented import path when completeness is uncertain. A compound assembled from all returned roots resolves the demonstrated loader defect; it does not certify arbitrary STEP product structures, external references or PMI.

Run `src/screen_mesh_v08b.py` against the retained STEP/GLB pair. Its three outcomes are sampled pass, observed-error rejection and inconclusive measurement. Preserve the witness points, query counts, resource failures and mesh diagnostics with the decision. Never interpret rejection as repaired geometry, or inconclusive measurement as acceptance. Compare alternative mesh generation separately, and freeze a changed policy before interpreting further files as unexposed validation.

The following files connect the report to the local supplement:

- `results/summary-v08b.json`: amended development and final-cohort classifications, distances and numerical cross-checks.
- `results/full-import-regression-v08b.json`: the partial and complete w20 source-import regression.
- `results/assembly-summary-v08b.json`: zero eligible final multi-root cases; `results/analytic-import-v08b`: separate synthetic controls.
- `experiments/local-v08b/protocol.json`: frozen implementation and source identities; `REPRODUCTION-v08.txt`: staging and replay instructions.

These are supplement-relative paths, not public download links. The report record is identified by DOI 10.5281/zenodo.23554447; no separate public evidence repository is provided with this release.

9 Reproducibility and declarations

The authoritative source remains `manuscript/working-paper.tex`, revised in place. The v0.7 source is retained under `experiments/local-v08`. Initial and amended protocols remain under `experiments/local-v08` and `experiments/local-v08b`. Frozen policy, source attribution and input hashes accompany `results/summary-v08b.json`. Separate development, held-out, exploratory and witness-check directories preserve chronology and failure accounting. `REPRODUCTION-v08.txt` describes the portable replay and standalone amended `screen_mesh_v08b.py` command. The initial v08 screen is superseded because it could evaluate an incomplete source. Earlier v0.5–v0.7 results and archives remain separate historical evidence.

Runs use Python 3.12.14, CadQuery 2.6.1, cadquery-ocp 7.8.1.1.post1, NumPy 2.5.3, Trimesh 5.1.1 and Gmsh 4.15.2. The original isolated Gmsh target installation omitted a data library; its unchanged library was extracted from the downloaded wheel. Wheel/library hashes and the vendor license are recorded. Gmsh binaries are not redistributed in the research package. Earlier SciPy-based analytic helicoid checks [14] remain historical supporting evidence, not verification of every new trimmed surface.

Prepared for BYBOLT, which has a commercial interest in fasteners and its website. Shiya Yuan is the sole author, affiliated with BYBOLT, as confirmed by the author. AI assistance supported implementation, analysis, literature screening and writing; final responsibility rests with the named human author. The author confirms that this work received no external dedicated funding and declares no additional competing interests beyond the BYBOLT affiliation and commercial interest disclosed above. External CAD retains its original CC BY 3.0 attribution and license; software dependencies retain their own terms. The author has confirmed the right to authorize distribution of the original report content. This report

and its corresponding LaTeX source are licensed under Creative Commons Attribution 4.0 International (CC BY 4.0): https://creativecommons.org/licenses/by/4.0/. The license does not cover the separate, unpublished research code, CAD models or measurement package, or replace any third-party rights. Numerical findings derive from retained executions, including incomplete and negative outcomes. No customer drawings or production data are used. Compilation and internal checks are not peer review.

Release preparation retains the original v0.8 evidence archive unchanged. A separately marked v0.9 candidate can normalize local operator path prefixes in textual logs and metadata, with original/candidate hashes and an explicit transformation record; numerical values, frozen code and CAD/mesh bytes are preserved. Embedded upstream source metadata is retained with the original source bytes. The separate evidence package remains local; public distribution of its original code and measurement material is not authorized by the article license. The file paths above document retained evidence and do not imply that a public supplement is currently available.

10 Conclusion

Local geometric probes detect all three previously exposed small-gate misses, and the screen preserves an unchanged accepted copy only after completed checks pass. Of twelve new files, eleven are measurable and one is empty. The sole observed baseline violation is detected by all three sampling methods. Independent comparison also exposes and helps correct incomplete source import, but the final cohort does not exercise unseen multi-root assemblies. The supported contribution is bounded detection and source-completeness evidence with controls, not a continuous fidelity guarantee or a general error-repair claim.

References

[1] FreeCAD community parts-library contributors. FreeCAD-library, pinned revision f5ca01cdcaa2f69b9295a6382719c2d90628780c. Source parts and per-file contribution history: https://github.com/FreeCAD/FreeCAD-library/tree/f5ca01cdcaa2f69b9295a6382719c2d90628780c. Asset license: https://github.com/FreeCAD/FreeCAD-library/blob/f5ca01cdcaa2f69b9295a6382719c2d90628780c/LICENSE-Assets. CC BY 3.0. Exact file paths and attribution accompany the supplement.

[2] X. Q. Zhang. Research on the Parametric Designing for Thread Fasteners. Applied Mechanics and Materials 80–81, 820–824, 2011. https://doi.org/10.4028/www.scientific.net/AMM.80-81.820. Abstract assessed.

[3] S. Bozherikov. Parametric modelling of bolted assembly and its components for industrial design and application. Journal of Applied Engineering Science 24(1), 188–202, 2026. https://doi.org/10.5937/jaes0-58641.

[4] Gumyr and contributors. cq_warehouse: parametric threaded fasteners. Software source. https://github.com/gumyr/cq_warehouse/blob/main/src/cq_warehouse/fastener.py.

[5] C. Gonzalez-Lluch et al. A survey on 3D CAD model quality assurance and testing tools. Computer-Aided Design 83, 64–79, 2017. https://doi.org/10.1016/j.cad.2016.10.003.

[6] CAx Interoperability Forum. Recommended Practices for Geometric and Assembly Validation Properties, v4.6, 2023. https://www.mbx-if.org/home/wp-content/uploads/2024/05/rec_prac_gvp_v46.pdf.

[7] D. Sap and D. P. Szabo. Representation and system-agnostic automated CAD interoperability testing based on model interchangeability. Journal of Computer Languages 85, 101365, 2025. https://doi.org/10.1016/j.cola.2025.101365. Publisher-indexed abstract/introduction assessed; full-text comparison pending.

[8] D. Sap and D. P. Szabo. An Automated Approach for the Discovery of Interoperability. arXiv:2001.10585v1, 2020. https://arxiv.org/html/2001.10585v1. Sections 2–4 assessed; not treated as the full text of the 2025 journal article.

[9] G. Li, Z. Luo, Y. Feng, L. Zhang, Y. Gai and N. Lei. Robust tessellation of CAD models without self-intersections. Journal of Computational Design and Engineering 13(1), 239–258, 2026; first published December 2025. https://doi.org/10.1093/jcde/qwaf134.

[10] Y. Zhou, D. Zint, N. Izadyar, M. Tao, D. Panozzo and T. Schneider. Topology-First B-Rep Meshing. arXiv preprint 2604.02141, 2026. https://arxiv.org/html/2604.02141v1. Preprint status; not treated here as verified journal acceptance.

[11] CadQuery contributors. Importing and Exporting Files. https://cadquery.readthedocs.io/en/stable/importexport.html.

[12] Khronos Group. glTF 2.0 Specification. https://registry.khronos.org/glTF/specs/2.0/glTF-2.0.html.

[13] Trimesh contributors. Proximity queries and source implementation. https://trimesh.org/trimesh.proximity.html. Local cross-check version 5.1.1.

[14] SciPy contributors. Bounded scalar minimization. https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize_scalar.html. Local numerical-oracle version 1.18.1.

[15] Gmsh contributors. Gmsh Reference Manual, version 4.15.2, 24 March 2026. Surface meshing, size controls, OpenCASCADE import and shape-distance API. https://gmsh.info/doc/texinfo/gmsh.html. Executed local version 4.15.2; OCC 7.8.1.

[16] C. Geuzaine and J.-F. Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. International Journal for Numerical Methods in Engineering 79(11), 1309–1331, 2009. https://doi.org/10.1002/nme.2579. Software attribution; current API settings are documented separately in the versioned manual.