PARTNEXT-MAT: LEARNING PIPELINE-AUTHORED MATERIAL PRIORS ON STRUCTURED 3D ASSETS

Anonymous authors
Paper under double-blind review

ABSTRACT

We introduce PartNeXt-Mat, a material-prior layer aligned with PartNeXt geometry, part masks, and hierarchy: 333,152 part assignments in 23,221 objects across 50 categories. Evidence fusion assigns a material, traceable grade resolution attaches applicable density and elastic priors, and verification-aware routing records solver caveats. A frozen 4,054-object visual benchmark studies part prediction and fixed-query completion. On its original split, category/name lookup reaches 82.9% accuracy, exposing near-total exact-key overlap. In an exploratory protocol withholding exact category/name keys from supervision, a fused visual-and-semantic model reaches 63.9% accuracy and 31.0% Macro-F1, compared with 57.8%/7.3% for a global prior, 54.5%/21.6% for train-only name/category backoff, and 57.8%/23.4% for a size-matched semantic MLP. The fused-minus-matched F1 gain is 7.60 points (key-cluster 95% interval 0.96–12.72). In name-disjoint completion, three observed labels raise set-attention accuracy by 1.67 points (0.70–2.63), although a support-aware table remains competitive. Numerical targets contain only 26 complete property tuples. The resource supports learning and completing authored, provenance-recorded part priors; it does not establish asset-specific measured mechanics.

1 INTRODUCTION

Making a textured 3D object simulation ready requires more than a mesh. A sofa may combine a rigid frame, compliant cushions, fabric covers, and metal legs; their density and elastic response determine which parts move, bend, or carry a load. Assigning plausible properties to each component matters for asset authoring, robotic interaction, and simulation, yet is difficult at scale: visual appearance is ambiguous, parts repeat, source measurements describe material specimens rather than the particular digital object, and a single object-level material misses mixed construction.

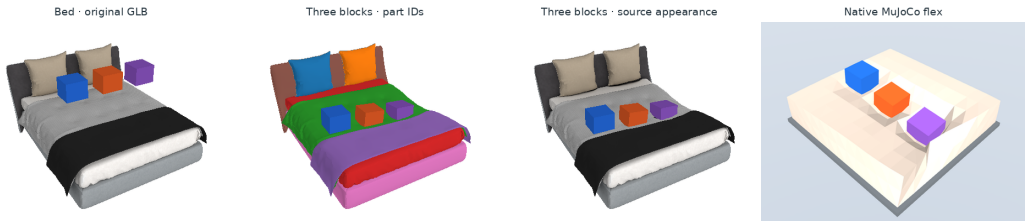

Figure 1: Illustrative simulator interface on an original held-out bed, not a learned-mechanics result. Three blocks with source-bucket densities 710/1200/2230 kg/m^3 fall on different soft regions. Their masses under an assumed scale are 1.99/3.36/6.25 kg; MuJoCo local flex peaks are 0.074/0.131/0.183 normalized units. Appendix M also fixes contact position to isolate mass. The original GLB is visual skinning and predicted priors do not set the solver.

Existing image-material datasets emphasize visible surface labels, while physical-property predictors often estimate fields or whole-object responses (Bell et al., 2015; Upchurch & Niu, 2022; Zhai et al., 2024; Dagli et al., 2026). Part/material datasets and CAD material recommendation establish important precedents (Ahmed et al., 2024; Comi et al., 2026; Bian et al., 2023). They do not, however,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

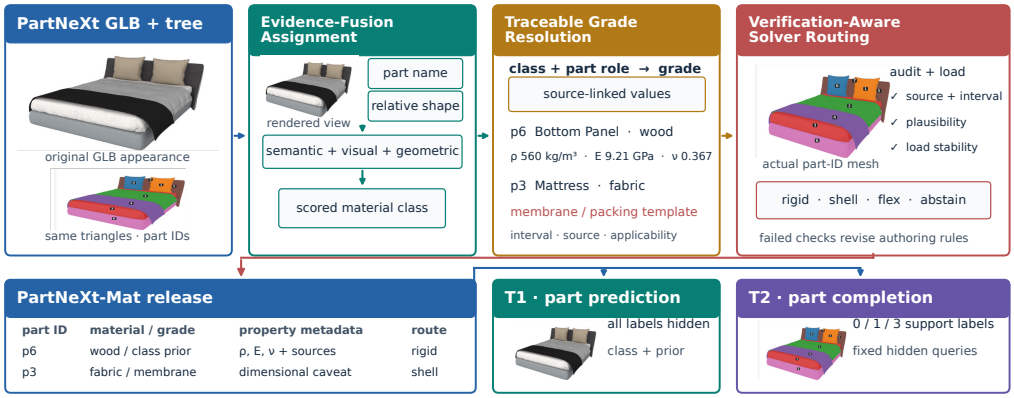

Figure 2: PartNeXt-Mat construction and downstream tasks. Every depicted object is an actual render of the same held-out PartNeXt Bed GLB (model `43781675a0`); the coloured view uses its original triangles and part masks. The two shown material rows are released assignments for that asset, not measured mechanics: part 6 has a wood prior, while part 3's fabric template is a membrane/packing assignment rather than a volumetric (ρ, E) target. Part names, rendered appearance, geometry, and head evidence propose classes; grade resolution attaches values, intervals, provenance, and applicability. Verification screens the release and routes parts to rigid, shell, volumetric, or abstention states, with findings revising authoring rules. The release supplies targets separately to prediction (T1) and fixed-query completion (T2); neither downstream model authors labels or calibrates a solver.

provide a direct test of whether a few known *part* materials help complete an aligned, source-traceable property layer over a large hierarchical mesh collection. A further difficulty is evaluation: when labels are authored using part names, high test agreement can reflect an exact semantic key rather than new visual or structural evidence.

We introduce **PartNeXt-Mat**, a material-prior layer for PartNeXt's existing geometry, masks, and semantic trees (Wang et al., 2025). Our core idea is to separate *material identity*, *source-linked numerical prior*, and *simulation applicability* instead of assigning an opaque physical triple to every mesh. This makes the release useful for part-level authoring while exposing the weaker evidence behind a reference fallback or a membrane parameter. It also makes the learning task testable: classifying a finite grade and looking up its properties is a strong baseline when the release uses only 26 complete class/property tuples.

The construction has three connected modules (Figure 2). **Evidence-Fusion Assignment** combines semantic, rendered-appearance, relative-geometry, and discriminative evidence to propose a part material. **Traceable Grade Resolution** maps that material and the part's role to a grade with per-quantity values, intervals, applicability, and provenance. **Verification-Aware Solver Routing** checks internal consistency and material plausibility, evaluates idealized load cases, records caveats, and routes rigid, shell, or volumetric parts; findings revise rules before the release is regenerated. These modules produce 333,152 part records in 23,221 objects across 50 categories. They produce *priors*, not direct measurements of each modeled object's mechanics.

We evaluate two downstream uses: unconditional part prediction and completion from zero, one, or three revealed coarse labels with a fixed hidden query set. Name-disjoint support/query episodes, semantic lookup, support-aware tables, matched context models, and a pretrained comparison prevent an easy copy baseline from being mistaken for structural reasoning. Three contributions follow: (i) a part-aligned, provenance- and applicability-aware material layer and explicit construction/verification pipeline; (ii) fixed-query protocols and controls that separate semantic repetition from part-specific and cross-part evidence; and (iii) reproducible experiments showing that a size-controlled fused input/model regime predicts authored assignments better on a frozen exact-key-supervision hold-out, and that observed labels help in narrower settings. A support-aware semantic table remains competitive in completion, while missing asset-specific measurements limit the physical claim.

2 RELATED WORK

2.1 MATERIAL DATASETS AND PART-ALIGNED ANNOTATION

Image material recognition has established both patch-level and dense segmentation benchmarks, including MINC and DMS (Bell et al., 2015; Upchurch & Niu, 2022). Such labels capture visible surface identity, which is valuable for appearance understanding but does not attach density or elastic parameters to the geometric components that a simulator consumes. 3DCoMPaT and 3DCoMPaT200 bring part and material labels into compositional 3D understanding (Li et al., 2022; Ahmed et al., 2024); thus part-material annotation itself is not new. Gen2Physics uses a large collection of material-annotated renderings from relabeled indoor assets and tests segmentation on manually labeled PartNet-Material views (Comi et al., 2026). Its component decomposition also supplies density and internal-structure choices for a generation pipeline. PartNet established fine-grained hierarchical part labels; PartNeXt provides our underlying meshes, part masks, and semantic hierarchy (Mo et al., 2019; Wang et al., 2025). Both target part understanding rather than material mechanics. Material-science tables report properties of specimens and named grades, not measurements of each digital object. Our release connects these interfaces: a part ID aligns a mesh region and hierarchy node to a material class, a grade, per-quantity provenance, and a solver-applicability flag. The corpus registry has broader grade coverage than the exported objects; the latter still use a small set of numerical templates, which we expose rather than hide. Relative to earlier datasets, our evaluation therefore asks whether source-linked priors can be predicted or conditionally completed on the *same aligned parts*, while explicitly checking semantic-key interpolation and refraining from treating an assigned grade as measured asset truth.

2.2 MATERIAL RECOMMENDATION AND CONTEXTUAL LEARNING

CAD systems often have assembly structure and part names before material specification. HG-CAD exploits body/assembly graphs to recommend a target material when the remaining component labels are provided (Bian et al., 2023). Deep Sets and Set Transformer establish permutation-invariant pooling and attention over element collections (Zaheer et al., 2017; Lee et al., 2019); Graphormer adds structural relation biases to attention (Ying et al., 2021). These methods motivate models that exchange information among parts, but a gain from context can be trivial if identically named components share the same authored label. We therefore compare an independent predictor, whole-object pooling, fixed hierarchy aggregation, set attention, and relation-biased attention using matched descriptors and approximately matched parameter counts. Even the independent predictor receives the whole-object image descriptor, so the comparison isolates interaction among part representations rather than simple access to object appearance. Our completion protocol differs from leave-one-out recommendation: it reveals zero, one, or three *coarse* labels, scores an unchanged hidden query set, and has a separate setting in which support and query part-name groups do not overlap. Same-name copying and a train-only support-aware co-occurrence table are baselines, not straw men; the table is competitive with the learned endpoint. Thus this work is related to graph-based recommendation, but its main contribution is a controlled measurement of what sparse observed labels add beyond semantic priors, rather than a claim that a new attention layer dominates existing CAD architectures.

2.3 VISUAL PHYSICAL INFERENCE AND SIMULATION-READY ASSETS

NeRF2Physics uses language-guided proposals and spatial features to infer physical properties, with evaluation involving object mass and measured pointwise properties (Zhai et al., 2024). VoMP predicts volumetric density, Young's modulus, and Poisson-ratio fields from visual geometry (Dagli et al., 2026). PhysGaussian couples 3D Gaussian rendering to MPM dynamics, while PhysDreamer distills interaction priors from video generation (Xie et al., 2024; Zhang et al., 2024). Other systems connect physical attributes with 3D generation or address rigid, deformable, and articulated assets, including PhysX-3D, PhysX-Omni, and UniPhysGen (Cao et al., 2025; 2026; Li et al., 2026). These works highlight the importance of a simulator-facing representation, but their outputs, supervision, and geometric units differ from our released part-level grade priors. In particular, a convincing rendered deformation does not establish that a network recovered a real object's modulus: scale, fill, attachments, mesh discretization, contact law, and boundary conditions also control the response. Our paper therefore separates its supervised benchmark from illustrative collision scenes. The benchmark tests agreement with released material assignments and applicable numerical priors,

compares prototype lookup with direct regression, and reports pretrained transfer on a matched part population where possible. The supplementary MuJoCo scenes use recorded solver states and original asset geometry for visualization, but their collision proxies and soft-body settings are authored rather than inferred by the benchmark model. Relative to physical-field prediction and simulation-ready generation, our contribution is an aligned, provenance-aware part interface and a deliberately bounded test of material-prior learning, not a claim of measured mechanics or end-to-end generated constitutive behavior.

3 PART-ALIGNED MATERIAL PRIORS

What is new. PartNeXt-Mat extends the existing PartNeXt geometry, part segmentation, and semantic hierarchy with material classes and source-linked physical-property assignments. The geometry and hierarchy are inherited from PartNeXt; our resource contribution is the aligned material layer and its evaluation protocols. Each record is keyed by an object identifier and a leaf-mask identifier. A separate storage identifier locates the GLB file and is never interpreted as an object category. A part mask indexes triangles within the corresponding GLB geometry. This alignment permits the same material annotation to supervise a rendered region, a geometric part, and a node in the object hierarchy.

The released table contains 333,152 pipeline-authored part assignments from 23,221 objects spanning 50 object categories. Ten material classes occur: metal, plastic, wood, fabric, glass, rubber, ceramic, foam, leather, and composite. There are no duplicate object–part keys. The release covers 333,152 of the 350,145 parts in the upstream annotation table (95.15%). Missing records denote unannotated parts, not a background material class. We preserve this distinction during rendering: unannotated triangles remain visible context but do not contribute supervised part targets.

Annotation construction and interpretation. The source corpus groups material-science records by named grade and tracks density, modulus, and Poisson-ratio evidence separately, with source identifiers and intervals for each quantity. The assignment pipeline first combines a category-conditioned part-name prior with rendered appearance, geometric evidence, and a discriminative material head. It caps the head's influence on names outside its training distribution. A second stage resolves a grade using part-role and structural constraints, then emits the material class, property values, interval semantics, provenance, and solver caveats. Appendix D describes these stages and the distinction between corpus coverage and the subset used in released labels. The purpose is to attach a plausible material prior to a digital asset. A source may contain measured material data, but assigning that source to an asset does not constitute a measurement of the asset itself. Available records do not establish disjoint training of the authoring head from this benchmark; our scores evaluate agreement with the released assignments, not independent material truth.

This construction makes semantic shortcuts a central experimental concern. A network that reproduces a category–part-name rule may accurately imitate the annotation pipeline while extracting little additional visual evidence. We therefore separate semantic, visual/geometric, and fused input protocols, and include explicit lookup baselines. Previously reported agreement with weak labels in the release documentation is a consistency check against authored rules; we do not reinterpret it as accuracy on an independently measured physical test set. The experiments below address learnability and conditional completion of the released priors.

Verification and its evidence limit. The supplied author protocol specifies eight internal audit suites, a 457-entry admissibility table, and mode-specific load screens for simulator-facing values. We independently recounted the released CSV: all 333,152 model–part keys are unique, every nonmissing point lies inside its own interval, 329,451 rows have a complete triple, and the solver-route totals sum to the release. The authors additionally report a random manual plausibility inspection of about 5% of final-release parts: they report viewing each selected part's mesh or render and assessing its material class and assigned property prior, with a 99.7% reported pass rate. The item-level sample, rubric, decisions, and corrections were not retained; we cannot reproduce that rate, compute a classwise error estimate, or use the sample as an independent model test. The authoring protocol's reported zero failures on internal checks likewise establish consistency with its rules, not ground truth for each part. Appendix E separates these evidence types.

Table 1: Evidence for release integrity and label validity. Only the first row is independently recountable from the available CSV; none measures an asset's physical response.

Evidence	Scope and reproducibility
CSV recount	333,152 unique keys, 26 complete templates, interval containment and derived route totals; script and report included.
Author rule screens	Eight suites, admissibility and idealized loads; input tables and execution logs unavailable.
Author manual review	About 5% sampled, 99.7% plausibility pass reported; no retained item IDs, rubric, or decisions.

Property templates and missingness. Density ρ, Young's modulus E, and Poisson's ratio ν are jointly present for 329,451 parts (98.89%). These complete rows contain only 26 distinct numerical triples. All 41,177 fabric rows share one membrane/packing template, so this release cannot benchmark fabric-subtype property inference. Consequently, predicting three scalars is not equivalent to learning three independent continuous physical processes. Our principal numerical baseline classifies the released fine property prototype and retrieves its values. We also evaluate a continuous regression head under the same features and training split.

The release distinguishes 247,652 bucket-level interquartile assignments, 81,213 class-level 5th–95th percentile assignments, and 586 reference-range assignments. Another 3,701 rows abstain from numerical assignment while retaining a material class. Interval types have different statistical meanings; we do not pool them into a calibrated predictive-confidence target. Provenance is retained for analysis and excluded from unknown-part inputs.

Dimensional applicability. The volumetric-property evaluation excludes density and modulus targets for 41,177 fabric rows flagged as membrane-modulus and packing-density assignments. In particular, a membrane stiffness cannot be converted to a volumetric Young's modulus merely because the CSV column is named in GPa. Fabric remains in material classification and in the valid Poisson-ratio evaluation. Foam values are interpreted as effective properties. Missing numerical targets are masked per quantity, and every numerical metric reports its eligible population. For valid volumetric entries, modulus is converted from GPa to Pa before taking logarithms.

Reproducible experimental subset. The visual experiment selects up to 100 objects per category by a fixed hash order, after requiring at most 256 annotated parts and a GLB size of at most 25 MiB. These thresholds are set before model fitting. This yields 4,059 candidate objects and 57,537 annotated parts, requiring 6.21 GiB of mesh downloads. Every candidate joins to the pinned upstream masks. Rendering failures are recorded by object and category; the successful-asset manifest is shared by all compared methods. We group identical GLB hashes and identical normalized, quantized vertex sets before assigning approximately 70/15/15% training/validation/test splits within categories. This detects exact and some normalized geometric duplicates, but does not guarantee removal of all approximate asset variants. The final successful-asset counts and exclusions are reported with the results.

Category generalization. A second protocol holds out Chair, Knife, Mug, Laptop Computer, and Skateboard categories. Any duplicate group touching a held-out category is assigned entirely to the held-out side. Validation uses only the remaining categories. This tests category transfer within the released resource; it is not an external measured-physics benchmark.

4 Construction and learning tasks

Problem definition. For an object o, PartNeXt supplies a mesh, an ordered part set $\mathcal{P}_o$, and a semantic tree T_o. Part $i \in \mathcal{P}_o$ has category and part-name metadata n_i, rendered appearance v_i, and relative geometry b_i. We seek a released record $A_i = (m_i, g_i, \{y_{iq}, I_{iq}, p_{iq}, a_{iq}\}_{q\in\{\rho,E,\nu\}}, r_i)$: material class m_i, grade g_i, each property's value y_{iq}, interval I_{iq}, provenance p_{iq}, applicability flag a_{iq}, and solver route r_i. A missing value is a recorded abstention, not a zero. The source corpus

describes material grades; assigning one to a mesh part is an inference, not a measurement of that part.

Pipeline overview. Figure 2 shows three connected modules. *Evidence-Fusion Assignment* proposes m_i from name, view, shape, and a discriminative head. *Traceable Grade Resolution* chooses g_i and emits property values with source and applicability metadata. *Verification-Aware Solver Routing* checks release-level invariants and plausible loading, attaches caveats, and chooses a compatible simulator family. Findings revise the rules and trigger full regeneration. The resulting table provides targets for unconditional prediction (T1) and sparse-label conditional completion (T2); these downstream networks do not author the released labels.

4.1 Evidence-Fusion Assignment

Let $\mathcal{C}$ be the material classes and $\pi_i(m)$ the category/name-conditioned semantic prior for $m \in \mathcal{C}$. The channel set $\mathcal{K} = \{\text{vis}, \text{geo}, \text{head}\}$ provides log-evidence $\ell_{ik}(m)$ for part i and class m. The authoring rule has the log-linear form

$$s_i(m) = \log \pi_i(m) + \sum_{k \in \mathcal{K}} \lambda_{ik} \min\{\ell_{ik}(m), \kappa_{ik}\}, \qquad \hat{m}_i = \arg\max_{m \in \mathcal{C}} s_i(m). \tag{1}$$

Here λ_{ik} weights an available channel and κ_{ik} caps its contribution. Visual colour is gated where the surface evidence is unreliable; the head's contribution is capped for out-of-distribution names. The semantic prior remains explicit because it is a powerful authoring signal and a potential benchmark shortcut. The supplied curation report's channel comparisons use weak internal labels, so we treat them as pipeline diagnostics rather than independent accuracy (Appendix D).

4.2 Traceable Grade Resolution

A deterministic rule R resolves a grade using the proposed class, part role, name, tree context, and grade registry $\mathcal{G}$:

$$g_i = R(\hat{m}_i, n_i, T_o; \mathcal{G}), \qquad (y_{iq}, I_{iq}, p_{iq}, a_{iq}) = Q_q(g_i), \quad q \in \{\rho, E, \nu\}. \tag{2}$$

Q_q retrieves the value, interval type, provenance, and dimensional applicability for quantity q; it returns a missing-value flag if no supported value exists. The registry preserves source-derived, reference-range, and class-imputed evidence separately for each quantity. Intervals can be grade interquartile ranges, class percentile ranges, or published reference ranges, and must not be interpreted as interchangeable uncertainty estimates. The released CSV has 26 complete class/property tuples and 3,701 rows that retain a class but abstain from a triple. Fabric's membrane-modulus and packing-density values are flagged as non-volumetric; foam modulus is effective rather than a cell-wall constant.

4.3 Verification-Aware Solver Routing

The third module evaluates a heterogeneous release-wide vector of checks $V(D)$ on the candidate table $D = \{A_i\}$ and assigns routes

$$r_i = R_{\text{sim}}(\hat{m}_i, \{y_{iq}, a_{iq}\}_q), \qquad V(D) = (V_{\text{key}}, V_{\text{interval}}, V_{\text{source}}, V_{\text{plaus}}, V_{\text{load}}, V_{\text{coupling}}). \tag{3}$$

The components test key/schema integrity, interval containment, provenance closure for corpus-backed values, name/material admissibility, idealized compression or bending and explicit-step stability, and mixed-object coupling. R_{sim} chooses rigid, cloth/shell, MPM elasticity, FEM elasticity, or abstention according to the recorded applicability and caveats. Author-reported findings update the rule set and regenerate the table before checks are rerun. The load cases screen implausible *priors*; they are not measured responses of the PartNeXt assets. The complete authoring rules and audit logs were not supplied, so these equations specify the interface rather than an end-to-end runnable release generator. Appendix E distinguishes reported checks from what is independently recountable in the CSV.

4.4 DOWNSTREAM PREDICTION AND COMPLETION

T1 predicts all annotated test parts from allowed rendered and geometric descriptors, with optional original category/name/tree semantics. T2 first fixes support candidates S and hidden queries Q by a target-independent seed; budgets $b \in \{0, 1, 3\}$ reveal nested coarse-class labels on S while scoring the same Q. A second protocol requires disjoint normalized part-name groups across S and Q. No released grade, property value, provenance, or hidden label enters an unknown part's input. All visual models use shared frozen DINOv2 features (Oquab et al., 2023); independent, pooled, fixed-graph, set-attention, and relation-biased heads have matched widths. The relation head adds parent, common-ancestor, and tree-distance biases to attention (Appendix C). A model predicts a finite grade prototype z_i and retrieves applicable (ρ, E, ν) from its public dictionary; a joint regression head is compared under the same feature and split protocol.

5 EXPERIMENTS

Data and protocol. The visual benchmark uses a frozen hash-selected subset of PartNeXt-Mat, with identical successful meshes, part masks, and grouped train/validation/test splits for every fitted method. It tests T1 prediction, ordinary fixed-query T2 completion, and a name-group-disjoint T2 variant; five entire object categories form a separate transfer split. A post-review exact category/name-key holdout masks supervision on disjoint key sets within the grouped object splits. Numerical metrics exclude inapplicable fabric density/modulus and always retain eligible queries even when a prediction needs fallback. Object-balanced material accuracy and ten-class Macro-F1 are primary. Physical-prior errors use absolute $\log_{10} \rho$, $\log_{10} E$, and raw ν differences. Geometry-group bootstrap intervals resample objects with all their parts and episodes; three fit seeds describe optimization variability.

The frozen visual benchmark contains 4,054 objects and 57,509 parts in 3,979 geometry groups, after excluding 5 candidates with unresolved preprocessing failures. Its training, validation, and ID-test sets contain 2,802, 618, and 634 objects. Ordinary completion evaluates 498 objects; disjoint-name completion evaluates 494. Each protocol preserves its own fixed queries across budgets. In the ID test, 659 of 8,783 annotated parts are invisible in the six rendered views and remain evaluated.

Implementation and comparisons. Six aligned RGB/part-mask renders per mesh supply frozen DINOv2-B patch features; each model also receives the same whole-object descriptor and 25 geometric/appearance summaries. Optional category, part-name, and tree-path features use a fixed 256-dimensional hash. Independent MLP, mean pool, fixed hierarchy, set attention, and relation attention use two residual blocks at approximately matched active parameter counts *within each input regime*; modality parameter counts differ. All three-seed fits use AdamW, batch size 32, learning rate 3×10^{-4}, weight decay 0.01, clipping at one, and a cosine schedule capped at 50 epochs with ten-epoch validation patience. Validation selects checkpoints and copying thresholds. Train-only global/category/name lookups, same-name and visual-support copying, and a support-aware semantic table are strong controls. Modality, hierarchy-bias, numerical-head, observation-budget, and equal-update data-fraction ablations test the role of evidence, structure, and supervision. VoMP and Qwen2.5-Omni are evaluated as pretrained transfers on explicitly matched populations, not as matched-backbone architecture tests. Appendices A–Q contain full protocols, transfer adapters, release audits, illustrative solver scenes, source bundles, and all held-out part-property tables.

5.1 SEMANTIC BASELINES REMAIN DIFFICULT TO BEAT

Table 2 shows that category and part names explain much of the released assignment. Exact training category/name keys cover 99.77% of ID-test parts; 99.00% have their test class and 97.68% their fine prototype somewhere under that key in training (Appendix G). The ID task is therefore largely exact-key interpolation. Lookup reaches 82.9% object-balanced accuracy; semantic-only learning reaches 67.7% Macro-F1, compared with 54.6% for visual/geometric learning. Fusion does not improve material classification over semantic-only prediction. Mean log-modulus error decreases from 0.384 to 0.364 dex, but the paired difference interval, $[-0.051, 0.009]$, includes zero. Poisson-ratio error decreases by 0.00146, with reduction interval $[0.00012, 0.00292]$; fusion's evidence is quantity-dependent.

Table 2: Unconditional prediction on the common grouped object split. S: category/name/path semantics; V: frozen visual plus geometry/appearance descriptors. Accuracy and F1 are object-balanced percentages (mean ± sample SD over three training seeds). Numerical columns show mean absolute prior error; full seed statistics accompany the code. Log errors are in dex; ν is dimensionless. All ten classes enter F1.

Predictor	Input	Accuracy ↑	F1 ↑	$\log \rho \downarrow$	$\log E \downarrow$	$\nu \downarrow$
Global lookup	S	47.8	6.5	0.472	1.271	0.047
Category lookup	S	65.7	30.3	0.278	0.750	0.033
Name lookup	S	77.1	60.6	0.200	0.509	0.023
Category–name lookup	S	82.9	65.0	0.156	0.390	0.018
Independent	S	82.5±0.3	67.7±1.1	0.158	0.384	0.018
Independent	V	74.4±0.4	54.6±1.6	0.208	0.536	0.024
Independent	V+S	81.8±0.1	65.4±1.1	0.156	0.364	0.016
Mean pooling	V+S	81.8±0.6	65.3±0.4	0.159	0.375	0.017
Fixed graph	V+S	82.1±0.6	65.3±1.5	0.155	0.369	0.016
Set attention	V+S	82.1±0.5	67.1±0.6	0.159	0.366	0.016
Relation bias	V+S	81.5±0.9	65.3±0.6	0.163	0.377	0.017

Train-defined ambiguity and label-free descriptor-rotation diagnostics in Appendix G localize this pattern. Fusion reduces log-modulus prior error on ambiguous category–name keys without improving their material F1 (Table 6). Rotating part visual and geometry descriptors within exact category–name groups increases fused-independent log-modulus error by 0.056 dex (95% interval $[0.033, 0.079]$), despite leaving semantics, object appearance, and fitted weights fixed (Table 8).

Context is not uniformly beneficial in unconditional prediction. Set attention attains 67.1% F1, compared with 65.4% for independent prediction and 65.3% for relation bias. The matched relation-minus-set F1 difference is -1.77 points (95% interval $[-3.54, -0.42]$). These results motivate evaluating what supplied material information changes, rather than treating a contextual architecture as sufficient evidence of structural reasoning.

5.2 Held-out exact-key supervision reveals a fused-model gain

To remove near-total exact-key supervision overlap, an exploratory post-review protocol hashes each category/lowercased-name key into disjoint supervised train, validation, and test sets. Objects retain their geometry-group split; only matching key/split parts enter loss or evaluation. Other parts remain unlabeled context: 77 of the 79 test keys appear there in training, so this is not unseen vocabulary. On 296 test objects and 1,557 queried parts, fused independent prediction reaches 63.9% accuracy and 31.0% ten-class F1 (Table 3). The global prior reaches 57.8%/7.3%; train-only name/category backoff reaches 54.5%/21.6%. Against backoff, the fused F1 gain is 9.46 points with a paired key-cluster interval $[1.84, 15.33]$; its accuracy interval $[-0.45, 18.74]$ crosses zero. A size-matched semantic MLP (499,755 versus 502,267 parameters) reaches 57.8%/23.4%; fused-minus-matched F1 is 7.60 points with a key-cluster interval $[0.96, 12.72]$. Two further fixed hash salts give fused-minus-matched F1 point gains of 12.53 and 4.46 points (Appendix G). These exploratory splits have sparse rare classes; the key interval conditions on the first split and three fits. The results measure authored assignments, not independent material correctness or the causal effect of a particular visual cue.

5.3 Completion gains require a copying control

On ordinary episodes, same-name copying reaches 69.0% F1 and gains 2.9 points from three labels, exceeding the matched fused-input learned models in final F1 (Table 4). The semantic-only relation model is stronger still, reaching 70.8% F1 with a 3.2-point gain. Thus copying is competitive, but does not explain away all learned completion performance. A benchmark that omits either reference can misattribute benefits to visual or structural reasoning. The independent model remains exactly invariant to the observation budget, as expected when observed nodes cannot affect hidden nodes.

Disjoint-name episodes remove direct copying. With three labels, fused set attention improves from 65.9% to 67.1% F1 and from 80.6% to 82.3% accuracy. Its paired accuracy gain is 1.67 points (95% interval $[0.70, 2.63]$), whereas its 1.21-point F1 gain has interval $[-0.02, 2.69]$. Relation bias

Table 3: Exploratory exact-key supervision holdout: 296 test objects, 1,557 queried parts, 79 keys. Accuracy and ten-class F1 are object-balanced percentages; fitted rows give mean ± SD over three seeds. No exact test key enters the lookup, so it uses name/category/global backoff. Foam has no test support.

Predictor	Input	Accuracy ↑	F1 ↑
Global prior	–	57.8	7.3
Category prior	S	46.0	15.4
Name/category backoff	S	54.5	21.6
Independent	S	57.2±2.5	24.7±1.6
Size-matched MLP	S	57.8±1.7	23.4±0.8
Independent	V	61.3±1.6	28.2±1.2
Independent	V+S	63.9±0.8	31.0±1.9
Set attention	V+S	61.2±1.0	29.3±0.8
Relation bias	V+S	62.3±0.6	31.7±2.0

Table 4: Completion with three observed coarse labels. All rows use fused descriptors except the semantic-only (S), visual/geometric (V), and train-only support-aware semantic lookup controls. Object-balanced ten-class F1 and gains are percentages. Learned/copy rows show mean ± SD over three fitted seeds; the deterministic lookup pools three episode seeds. Ordinary and disjoint-name populations differ; queries stay fixed across budgets.

Predictor	Ordinary F1	Δ_{3-0}	Disjoint-name F1	Δ_{3-0}
Independent	66.1±1.2	0.0±0.0	65.7±0.7	0.0±0.0
Mean pooling	66.2±0.9	0.8±2.1	65.0±0.8	0.0±2.0
Fixed graph	67.1±1.3	2.2±0.3	65.5±1.3	0.9±1.0
Set attention	67.1±1.3	2.0±0.8	67.1±0.7	1.2±0.4
Relation bias	65.9±1.4	1.4±0.9	65.6±0.5	1.1±0.9
Relation bias (S)	70.8±0.2	3.2±1.1	68.0±0.5	0.8±1.2
Relation bias (V)	57.9±1.4	0.9±0.2	56.5±0.7	0.6±0.1
Same-name copying	69.0±1.2	2.9±0.3	65.7±0.7	0.0±0.0
Visual copying	65.7±1.7	-0.4±0.7	65.7±0.7	0.0±0.0
Support-aware lookup	67.8	3.8	66.0	0.2

gains 1.07 F1 points ($[0.35, 1.85]$) but ends at 65.6%, below semantic-only relation bias at 68.0%. A train-only support-aware semantic table reaches 83.8% accuracy and 66.0% F1 on the same queries, versus 82.3% and 67.1% for fused set attention. Set-minus-table endpoint differences are -1.53 accuracy points ($[-3.16, 0.04]$) and $+1.07$ F1 points ($[-1.05, 3.09]$). An exploratory ambiguous-key analysis likewise finds no established endpoint advantage (Appendix G, Table 7). Supplied labels help some within-model outcomes without establishing a learned-context, hierarchy, or visual-fusion advantage.

Transfer and target controls. Prototype lookup reduces ID log-density/log-modulus error from 0.215/0.518 to 0.156/0.364 dex relative to joint regression. On five held-out categories, name lookup reaches 72.5% accuracy and 46.0% ten-class F1, above learned fusion. These controls limit the exact-key result to interpolation among authored priors; full ablations and pretrained-transfer results appear in Appendix G.

6 Conclusion

PartNeXt-Mat aligns structured 3D parts with traceable material and solver-applicable property priors. Across three exploratory exact-key splits, fusion improves ten-class F1 over a size-matched semantic model, though accuracy gains are uncertain. Observed labels can aid completion; finite templates, rare classes, and competitive semantic controls limit interpretation. Provenance and applicability fields support solver routing or abstention. The collision gallery illustrates authored settings, not validated material inference. Retained item-level judgments and measured force–displacement tests remain necessary for physical validation.

AI USE STATEMENT

Generative AI tools were used solely for grammar correction and minor language polishing of text written by the authors. No generative AI tools were used for any research tasks requiring disclosure under the ICLR 2027 AI Policy. All AI-assisted edits were reviewed by the authors, who take full responsibility for the final content of this work.

ETHICS STATEMENT

We reuse PartNeXt geometry and parts to study authored material priors. Incorrect assignments could mislead simulations; we distinguish applicable targets and template fallbacks from measured mechanics. Upstream licenses remain applicable.

REPRODUCIBILITY STATEMENT

The artifact records source revisions, checksums, included and excluded assets, grouped splits, model configurations, and per-object predictions. Baselines share fixed observation episodes and training-only normalization; Appendix A gives full protocols.

REFERENCES

Simeon Adebola, Chung Min Kim, Justin Kerr, Shuangyu Xie, Prithvi Akella, Jose Luis Susa Rincon, Eugen Solowjow, and Ken Goldberg. Botany-Bot: Digital Twin Monitoring of Occluded and Underleaf Plant Structures with Gaussian Splats. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1839–1846, 2025. doi: 10.1109/IROS60139.2025.11247474. URL https://huggingface.co/datasets/SimeonOA/Botany-Bot.

Mahmoud Ahmed, Xiang Li, Arpit Prajapati, and Mohamed Elhoseiny. 3DCoMPaT200: Language grounded large-scale 3D vision dataset for compositional recognition. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/f4f094d1c87437fb9157a9eae5a3def0-Abstract-Datasets_and_Benchmarks_Track.html.

Sean Bell, Paul Upchurch, Noah Snavely, and Kavita Bala. Material recognition in the wild with the materials in context database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3479–3487, 2015. URL https://openaccess.thecvf.com/content_cvpr_2015/html/Bell_Material_Recognition_in_2015_CVPR_paper.html.

Shijie Bian, Daniele Grandi, Tianyang Liu, Pradeep Kumar Jayaraman, Karl Willis, Elliot Sadler, Bodia Borijin, Thomas Lu, Richard Otis, Nhut Ho, and Bingbing Li. HG-CAD: Hierarchical graph learning for material prediction and recommendation in computer-aided design. *Journal of Computing and Information Science in Engineering*, 24(1):011007, October 2023. ISSN 1944-7078. doi: 10.1115/1.4063226. URL http://dx.doi.org/10.1115/1.4063226.

Ziang Cao, Zhaoxi Chen, Liang Pan, and Ziwei Liu. PhysX-3D: Physical-Grounded 3D Asset Generation. In *Advances in Neural Information Processing Systems*, 2025. URL https://arxiv.org/abs/2507.12465.

Ziang Cao, Yinghao Liu, Haitian Li, Runmao Yao, Fangzhou Hong, Zhaoxi Chen, Liang Pan, and Ziwei Liu. PhysX-Omni: Unified Simulation-Ready Physical 3D Generation for Rigid, Deformable, and Articulated Objects. *arXiv preprint arXiv:2605.21572*, 2026. URL https://arxiv.org/abs/2605.21572.

Mauro Comi, Jordi Serrano Berbel, Kevis-Kokitsi Maninis, Philipp Henzler, and Manuel Sanchez. Gen2Physics: Grounding generated 3D meshes in physics via multi-view material decomposition. *arXiv preprint arXiv:2608.23869*, 2026. URL https://arxiv.org/abs/2608.23869.

Rishit Dagli, Donglai Xiang, Vismay Modi, Charles Loop, Clement Fuji Tsang, Anka He Chen, Anita Hu, Gavriel State, David I. W. Levin, and Maria Shugrina. VoMP: Predicting Volumetric Mechanical Property Fields. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2510.22975.

Juho Lee, Yoonho Lee, Jungtaek Kim, Adam R. Kosiorek, Seungjin Choi, and Yee Whye Teh. Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. In *International Conference on Machine Learning*, 2019. URL https://arxiv.org/abs/1810.00825.

Xian Li, Rong Wei, Lujie Yang, Haolin Huang, Junyuan Fang, Siliang Tang, Jun Xiao, Rui Tang, and Juncheng Li. UniPhysGen: Unified Physical Grounding for Simulation-Ready 3D Assets. *arXiv preprint arXiv:2607.13586*, 2026. URL https://arxiv.org/abs/2607.13586.

Yuchen Li, Ujjwal Upadhyay, Habib Slim, Ahmed Abdelreheem, Arpit Prajapati, Suhail Pothigara, Peter Wonka, and Mohamed Elhoseiny. 3D CoMPaT: Composition of Materials on Parts of 3D Things. In *European Conference on Computer Vision*, 2022. URL https://www.ecva.net/papers/eccv_2022/papers_ECCV/html/3631_ECCV_2022_paper.php.

D. V. Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005, 1956. doi: 10.1214/aoms/1177728069.

Kaichun Mo, Shilin Zhu, Angel X. Chang, Li Yi, Subarna Tripathi, Leonidas J. Guibas, and Hao Su. PartNet: A Large-Scale Benchmark for Fine-Grained and Hierarchical Part-Level 3D Object Understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 909–918, 2019. URL https://openaccess.thecvf.com/content_CVPR_2019/html/Mo_PartNet_A_Large-Scale_Benchmark_for_Fine-Grained_and_Hierarchical_Part-Level_3D_CVPR_2019_paper.html.

Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning Robust Visual Features without Supervision. *arXiv preprint arXiv:2304.07193*, 2023. URL https://arxiv.org/abs/2304.07193.

SAM 3D Team, Xingyu Chen, Fu-Jen Chu, Pierre Gleize, Kevin J. Liang, Alexander Sax, Hao Tang, Weiyao Wang, Michelle Guo, Thibaut Hardin, Xiang Li, Aohan Lin, Jiawei Liu, Ziqi Ma, Anushka Sagar, Bowen Song, Xiaodong Wang, Jianing Yang, Bowen Zhang, Piotr Dollár, Georgia Gkioxari, Matt Feiszli, and Jitendra Malik. SAM 3D: 3Dfy Anything in Images. *arXiv preprint arXiv:2511.16624*, 2025. URL https://arxiv.org/abs/2511.16624.

Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018. URL https://arxiv.org/abs/1708.00489.

Paul Upchurch and Ransen Niu. A dense material segmentation dataset for indoor and outdoor scene parsing. In *European Conference on Computer Vision*, 2022. URL https://www.ecva.net/papers/eccv_2022/papers_ECCV/html/6243_ECCV_2022_paper.php.

Penghao Wang, Yiyang He, Xin Lv, Yukai Zhou, Lan Xu, Jingyi Yu, and Jiayuan Gu. PartNeXt: A Next-Generation Dataset for Fine-Grained and Hierarchical 3D Part Understanding. In *Advances in Neural Information Processing Systems*, 2025. URL https://arxiv.org/abs/2510.20155.

Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, and Chenfanfu Jiang. PhysGaussian: Physics-Integrated 3D Gaussians for Generative Dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. URL https://openaccess.thecvf.com/content/CVPR2024/papers/Xie_PhysGaussian_Physics-Integrated_3D_Gaussians_for_Generative_Dynamics_CVPR_2024_paper.pdf.

Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://papers.nips.cc/paper_files/paper/2021/hash/f1c1592588411002af340cbaedd6fc33-Abstract.html.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and Alexander Smola. Deep Sets. In *Advances in Neural Information Processing Systems*, 2017. URL https://arxiv.org/abs/1703.06114.

Albert J Zhai, Yuan Shen, Emily Y Chen, Gloria X Wang, Xinlei Wang, Sheng Wang, Kaiyu Guan, and Shenlong Wang. Physical property understanding from language-embedded feature fields. In *CVPR*, 2024.

Tianyuan Zhang, Hong-Xing Yu, Rundi Wu, Brandon Y. Feng, Changxi Zheng, Noah Snavely, Jiajun Wu, and William T. Freeman. PhysDreamer: Physics-Based Interaction with 3D Objects via Video Generation. In *European Conference on Computer Vision*, 2024. URL https://www.ecva.net/papers/eccv_2024/papers_ECCV/papers/00270.pdf.

A REPRODUCTION DETAILS

Reading guide. The first sections document the benchmark protocol, statistical tables, pretrained interfaces, and annotation pipeline. The later sections give illustrative solver scenes and per-part held-out asset galleries; the offline video site provides a searchable index and playable trajectories. The scenes are interface examples, not validation of the learned material predictions.

Pinned releases. The material annotations, part masks, and meshes use pinned source revisions recorded in the reproducibility artifact. All Arrow shards and each selected GLB are checked against expected file size and SHA256. Split manifests and preprocessing failures are saved before fitting the benchmark models.

Software environments. The core benchmark used Python 3.13.11, PyTorch 2.10.0, and the recorded core package versions; MuJoCo/Genesis examples used a separate Python 3.11.15 environment with MuJoCo 3.10.0 and Genesis 1.2.2. The artifact includes an environment matrix and separate simulator package pins. VoMP's upstream requirement list is reference material, not installed into the adapted core benchmark environment.

Geometry ordering. Mask indices refer to geometry storage order. In the installed trimesh version, scene-graph traversal can reverse this order. We preserve world transforms while sorting the dumped triangle meshes by their geometry insertion order, consistent with the release's geometry-verification script. Non-triangle scene primitives are omitted. Mesh count and per-mesh face count must agree with the annotation, every face index must be valid, and annotated masks must be disjoint. Assets with ambiguous instancing or unresolved mismatches are excluded and reported. We do not repair a mismatch by matching material labels or by guessing from part appearance.

Rendering and features. Meshes are centered and isotropically scaled so their longest extent is 1.5. Six perspective cameras surround the object at equally spaced azimuths, alternating two elevations. RGB and integer part-ID attachments share the same rasterization and depth buffer at 224×224 pixels. Base-color texture or vertex color is combined with a simple diffuse illumination term. The representation is appearance evidence, not a physically based measurement apparatus; transmission and complex reflectance are not reconstructed. DINOv2-B/14 remains frozen. Images are neither resized nor center-cropped after mask rendering, keeping its 16×16 patch grid aligned with the masks. A part's patch contribution is proportional to its pixel overlap, summed across views. Textures are resized to at most 1024×1024 for temporary rendering and are not saved as an additional dataset copy.

The 25 geometric/appearance features are center (3), axis-aligned extent (3), logarithmic surface area (1), logarithmic face count (1), vertex-covariance eigenvalues (3), mean and standard deviation of RGB (6), mean normals and mean absolute normals (6), and area-independent face-count-weighted metallic/roughness factors (2). Geometry uses normalized asset coordinates, so these descriptors do not establish real-world physical scale. Hash-based semantic features use a fixed signed 256-dimensional feature map. Object category, complete part name, complete ancestor path, and part-name word tokens are encoded separately.

Prototype construction. The fine vocabulary is the released set of material-class/property tuples, including the all-missing ceramic tuple. This dictionary is declared benchmark prior knowledge. No test-part assignment is exposed as an input. Density errors use $\log_{10}(\rho\,[\mathrm{kg/m^3}])$, modulus errors use $\log_{10}(E\,[\mathrm{Pa}])$, and Poisson-ratio errors use the dimensionless raw value. Numerical error is computed only where the reference has an applicable quantity, but an invalid predicted prototype never removes a query from the denominator. Its train-median fallback and fallback rate remain visible in the report.

Coarse classification maximizes summed prototype probabilities, whereas numerical lookup selects the single highest-probability prototype. These separate decisions can disagree on the coarse class; they are evaluated as distinct outputs and do not guarantee a coherent joint label/triple. The regression variant retains both classification losses and adds numerical Smooth-L_1 at weight one, scaling each quantity by its training standard deviation with floor 0.01. It uses the same material-classification validation criterion as lookup, rather than selecting a numerically optimal checkpoint.

Table 5: Three-label fixed-query completion. The support-aware table uses only training-set category/name prototypes and category-specific material-class co-occurrences, with two strengths selected on validation. It is deterministic and pools three observation seeds. Learned set attention uses fused inputs; its entries are mean ± SD over three fitted seeds. Populations differ between protocols.

Protocol	Predictor	Accuracy (%)	F1 (%)
Ordinary	Support-aware lookup	85.8	67.8
Ordinary	Set attention	83.7±0.1	67.1±1.3
Disjoint name	Support-aware lookup	83.8	66.0
Disjoint name	Set attention	82.3±0.6	67.1±0.7

Completion populations and aggregation. Ordinary episodes require at least four annotated parts. A seeded, target-independent permutation assigns three supports and a fixed complement of queries. Group-controlled episodes first remove digit tokens and normalize case and whitespace in part names. They reveal supports drawn from complete name groups and evaluate only names in disjoint groups. The maximal budget is always three observed coarse labels. Remaining unobserved support candidates are context nodes and are not added to the fixed query set at lower budgets.

Group-controlled eligibility can vary with the seeded group ordering. We attempt all three fixed seeds, retain eligible episodes, and average them within each object. Thus the grouped population includes objects with one, two, or three eligible episodes; it does not count repeated episodes as independent objects. Exact per-seed coverage accompanies the population manifest.

Support-aware semantic baseline. To test whether a learned context model is needed, we add a train-only conditional table. A category/name prototype prior backs off to name, category, and global counts; each training object contributes equal total weight. Within each training object, ordered material-class pairs give a category-specific co-occurrence likelihood, smoothed by five equivalent global-object pair tables. For a query prototype, the log prior receives the summed log likelihoods of the observed coarse classes with a tunable weight. A zero weight reduces to smoothed semantic lookup. Validation selects the prior smoothing from $\{0.1, 1, 10\}$ and co-occurrence weight from $\{0, 0.25, 0.5, 1, 2\}$ separately for ordinary and disjoint-name episodes, using observation seed zero and the mean accuracy/F1 criterion across budgets one and three. The table sees no rendered features or held-out target labels; test observations are exactly the coarse support labels available to the learned methods.

On disjoint-name episodes with three labels, this deterministic baseline reaches 83.8% accuracy and 66.0% ten-class F1, versus 82.3% and 67.1% for fused set attention. Paired set-minus-table intervals span zero for accuracy (-1.53 points, 95% interval $[-3.16, 0.04]$) and F1 ($+1.07$ points, $[-1.05, 3.09]$), averaged over the three fitted set models. The table's own three-versus-zero gain is $+1.31$ accuracy points ($[-0.22, 3.03]$) and $+0.23$ F1 points ($[-0.93, 1.44]$). Thus the learned model's within-model accuracy improvement from supplied labels remains an observed effect, but neither method has an established endpoint advantage over the other. Complete prediction records and paired geometry-group resamples accompany the artifact.

For each method, predictions are stored with object identifiers, query indices, observation seed, and budget. The primary object accuracy first averages query correctness within an episode, then averages episodes of the same object, and finally averages objects. Object-balanced Macro-F1 is computed from a confusion matrix in which each object's episodes share total weight one and each episode's queries share its weight uniformly. All ten released classes remain in the macro average, including a zero contribution when a class has no support. Conventional part-pooled metrics, class supports, and a secondary F1 over supported classes are reported separately. Numerical errors use the same nesting, restricting each quantity to valid reference targets. Confidence intervals resample geometry-duplicate groups, retaining their member objects and all observation episodes; parts and repeated observation seeds are not treated as independent samples.

Validation discipline. The maximum training schedule, optimizer, parameter-matching rule, and data partitions are shared across each matched comparison. Checkpoints are selected using validation-only object-balanced Macro-F1 and accuracy. Completion selection averages its three observation budgets at a fixed validation observation seed. Copy-baseline thresholds are fitted on validation

episodes. The test split is used for final reporting, never for selecting thresholds, epochs, widths, or material-word filters. Runtime verification includes fixed-query/nested-support assertions, hidden-target perturbation checks across all five model kinds, repeated-name group isolation, and invariance of object aggregation to duplicated observation episodes.

Data-fraction runs allocate equal optimizer-update budgets and retain validation-selected checkpoints, whose selected epochs can differ. Within-category fractions round upward and keep at least one training object per category; actual object counts accompany the run configurations. Reported bootstrap intervals are pointwise, not simultaneous across all comparisons. Resamples with no applicable object for a quantity are excluded for that quantity and their valid-draw counts are saved.

The released statistical archive additionally includes within-model observation-gain intervals for semantic-only and visual-only completion. These reporting analyses were added after aggregate results were inspected and are exploratory; they use the same bootstrap specification and do not change any fit or evaluation episode.

Parameter matching is performed separately for each descriptor, task, and head setting. Consequently, modality ablations change the input encoder's parameter count as well as its evidence; they are comparisons of input regimes, whereas the set/relation pair isolates a 24-parameter structural addition at identical width. The fixed material-word list uses standalone-word matching in names and paths; joined forms and other indirect semantic cues can remain.

During implementation, an initial copy-baseline pass reused ordinary-validation thresholds for grouped episodes. Following inspection, we corrected this population mismatch by selecting each protocol's threshold on its corresponding validation episodes, with the original grid and selection criterion unchanged. Final tables use the corrected evaluation. The original outputs and the dated amendment remain in the experiment record; this correction changes no learned checkpoint, split, or hidden query.

B PRETRAINED MODEL INTERFACES

VoMP rendering/backend adapter. We use the released geometry-model and material-VAE weights at model revision `f9639381b6b8012a8a975b5b9503aa2ea04f5709`, with upstream source pinned to `ac6826f25ca5cfa122eb7bacb74be91b5a9eadf4`. The configuration uses a 12-block, 768-wide, 12-head shifted-window transformer at voxel resolution 64 and window size eight, with absolute positional embeddings. A Torch attention backend preserves the released weights and window operations while avoiding optional sparse/flash dependencies. Strict parameter loading and numerical checks of attention and positional embeddings are retained in the implementation logs.

The released camera sampler supplies 150 seeded views at radius two and 40-degree field of view. We render 512×512 images against black with the same ModernGL material adapter used for preprocessing. This simple diffuse renderer differs from the original Blender/Cycles pipeline. DINOv2-L/14 with register tokens processes images resized to 518×518; its pre-final-normalization 37×37 patch features are projected to voxel locations and averaged over views, including the released float16 storage round trip. We retain at most 32,768 unique occupied voxel indices with a fixed seed. Features project from grid indices divided by 64 minus 0.5; prediction queries use voxel centers. Unique-index handling and rendering are declared adaptations.

For each annotated part, 512 deterministic-seed, area-uniform surface samples query the nearest predicted voxel. Coordinate-wise medians in $(\log_{10}\rho, \log_{10}E, \nu)$ yield part estimates. This aggregation neither estimates a volume-weighted mass nor guarantees that the resulting triple belongs to a constitutive-material manifold. We preselect the first 100 frozen test objects by the existing object hash before observing performance. Every compared in-domain method is evaluated on exactly the successful VoMP objects and their complete annotated part sets. Failures remain in the coverage report. No material label is passed to VoMP inference.

Qwen2.5-Omni-7B. The pretrained checkpoint is verified against its release checksums. Only the image/text Thinker is evaluated; speech generation is disabled. For each part, we choose the six-view render with the most visible target pixels. The prompt contains that full image with a red target boundary, a padded crop resized to 224×224, and the original category, part name, and hierarchy

path. Invisible parts receive a blank crop and an explicit visibility statement. The class options, in order, are ceramic, composite, fabric, foam, glass, leather, metal, plastic, rubber, and wood, mapped to A–J. The system prompt is the model's standard prompt. Exact user text and selected IDs are saved in the inference manifest.

Inference ranks the next-token logits for the ten single-token letters. A frozen head containing precisely those rows of the original language-model output matrix avoids computing unrelated vocabulary logits. It introduces no trained classifier parameters. Inputs use left padding, two images per example, and fixed 224^2 image-pixel limits; the final real prompt token is the scoring position. The resulting ten-way normalization is constrained answer ranking, not a calibrated physical-material posterior or unrestricted text generation. Option ordering is fixed before test inference.

The multimodal implementation check uses unequal-length real prompts and confirms image-token grids, finite normalized outputs, no missing Thinker weights, and close agreement between single-example and batched scoring. The formal run attempts all frozen ID-test objects. Material comparisons require identical dataset hashes, object IDs, and part ordering, and report visible/invisible coverage. Qwen is not assigned an invented continuous-property output. Pretraining overlap with the underlying assets is unknown.

Active-query computation. Given revealed supports S, fixed queries Q, and remaining candidates $C \setminus S$, look-ahead selects

$$j^* = \arg\min_{j \in C \setminus S} \sum_c p_\theta(c_j = c \mid x, c_S) \frac{1}{|Q|} \sum_{i \in Q} H[p_\theta(c_i \mid x, c_S, c_j = c)] \,. \tag{4}$$

The conditional network need not define a consistent joint posterior. This objective is expected model-predicted entropy, not a proved mutual-information criterion or a guarantee of improved accuracy. Information-based experimental design has a long history (Lindley, 1956), and geometric diversity is established in active learning (Sener & Savarese, 2018). The experiment evaluates simple acquisition rules with the learned part-completion model; it does not reproduce the full core-set algorithm.

Each active episode has a fixed candidate cap of eight, observation seed zero, and at least four candidates. It retains all other name groups as hidden queries. The first 100 eligible test objects are selected by an independent object hash, with the eligible and selected counts saved. The diversity policy starts with the highest-entropy candidate, then maximizes distance to observed candidates in the normalized descriptor space. Look-ahead enumerates each of the ten possible coarse labels for every remaining candidate, evaluates the conditional model in batches of 16, and minimizes expected mean query entropy. It uses at most $10(8+7+6) = 210$ hypothetical states for a three-label trajectory. We record candidate indices, query indices, selected labels, hypothetical-state counts, and approximate per-policy wall-clock cost. Timings come from a shared GPU server and do not establish isolated hardware throughput. One fixed random acquisition sequence per object is used; the three fitted models do not constitute three independent acquisition-seed trials.

C Downstream model and completion details

Fixed queries and label budget. For an object with at least four annotated parts, a target-independent seeded permutation chooses three support candidates S; the complement is the fixed query set Q. Budgets $b \in \{0, 1, 3\}$ reveal nested prefixes of S. Unrevealed candidates remain context nodes and are unscored. Consequently, zero-budget T2 and all-part T1 have different populations. The name-controlled version groups normalized part names first and requires support and query groups to be disjoint. Every method uses the same seeds 0, 1, and 2 and the same scored queries within an episode. This removes exact repeated-name copying but does not prove that nearby geometric variants are absent.

Features and context. Six RGB renders and aligned integer masks provide part-aligned DINOv2-B patch features, pooled by visible pixels. The mean class-token embedding gives every part the same whole-object descriptor. Twenty-five geometric and appearance summaries and optional original category/name/ancestor-path hashes complete the descriptor x_i. Invisible parts retain geometry and a zero visual vector. Normalization uses training objects only. Identifiers, released target grades,

property values, confidence, provenance, and source quantities are never used as unknown-part inputs. An observed class is embedded only at its revealed support node.

Independent residual MLPs, object mean pooling, fixed hierarchy aggregation, set attention, and relation-biased attention each have two residual blocks. The independent model receives the whole-object visual token but no support-label embedding. The relation head modifies set attention as

$$a_{ij}^{(h)} = \mathrm{softmax}_j\left(\frac{\langle W_q^{(h)} u_i, W_k^{(h)} u_j\rangle}{\sqrt{d_h}} + (b^{(h)})^\top r_{ij}\right). \tag{5}$$

Here u_i is the encoded part descriptor, $W_q^{(h)}$ and $W_k^{(h)}$ are head-h query/key projections, d_h is that head's width, and r_{ij} contains shared-parent, normalized common-ancestor-depth, and clipped tree-distance features. The vector $b^{(h)}$ contains the corresponding learned coefficients. The fixed graph instead normalizes sibling indicators and exponentially decaying tree-distance weights. These are semantic relationships, not physical-contact constraints.

Prediction and optimization. For prototype z_i in the released finite dictionary, each model predicts $p_i(z)$. The class posterior sums compatible prototype probabilities, $p_i(m) = \sum_{z:m(z)=m} p_i(z)$. A prototype cross-entropy plus 0.5 times class cross-entropy is averaged over supervised parts and then objects. Numerical lookup chooses the maximum-posterior prototype and retrieves valid properties; if the predicted prototype lacks an eligible quantity, the training-median fallback is applied and counted. The joint-regression comparison adds masked training-scale-normalized Smooth-L_1 error. All model widths, seeds, optimizer settings, and checkpoint selection follow Section 5. The released prototype dictionary is declared public task prior knowledge, not evidence that a network can infer arbitrary continuous constitutive parameters.

Active observation. Active completion keeps four to eight support candidates in name groups disjoint from fixed queries. Random, entropy, descriptor-diversity, and one-step model-entropy look-ahead use a frozen trained relation head. The true label is read only after a candidate is chosen. The model's conditional probabilities need not form a coherent joint posterior; look-ahead is a heuristic rather than a guaranteed information-gain optimizer. Detailed acquisition computation and coverage follow below in Appendix B.

D MATERIAL CORPUS AND ANNOTATION PIPELINE

The companion curation specification constructs a grade-level material table from openly reusable material-science records. A grade is identified by material name rather than by proximity to a desired density or modulus. For each quantity $q \in \{\rho, E, \nu\}$, the record carries a point value, interval, contributing source IDs, record count, and provenance. The provenance distinguishes source-derived values, published reference values, class imputations, and withheld values. These fields are not collapsed into a single claim that every quantity of every asset was measured. The source taxonomy admits more classes than the released object annotations: the exported CSV contains ten material classes, 23 nonempty assigned bucket identifiers, and 26 distinct complete class/property tuples. Its per-quantity fields contain 42 literal source-ID strings, representing 40 underlying IDs after three `proxy::` aliases are normalized. These are counts of the downstream release, not of every candidate grade or source in the broader curation registry.

Material assignment proceeds in two stages. First, category and part-name evidence supplies a semantic prior, combined with rendered surface appearance, relative geometry, and a discriminative material head. The supplied specification gates weak colour evidence and caps the head's influence on part names outside its training distribution. This design is relevant to evaluation because a model may learn the same category–name regularities that helped author its targets. Second, the pipeline resolves a material grade using recorded part-role, name, and structural constraints. It emits per-quantity intervals and provenance alongside (ρ, E, ν) and permits abstention when no supported numerical triple is available. The release's interval kinds distinguish grade-bucket interquartile ranges, class-level percentile ranges, and published reference ranges; they are not interchangeable confidence intervals.

The source curation documents contain weak-label channel diagnostics from unreconciled configurations: a supplied pipeline configuration records 92.2% weak-label agreement, while a separate

report states 92.8% and an earlier channel table states 95.9%. Their exact snapshots and evaluation populations cannot be matched from retained records. We therefore do not use these percentages as validation of the released labels or as a module ablation.

The exported simulation columns include derived elastic quantities and solver-routing metadata. Their meaning depends on material family. In particular, fabric rows use a membrane-modulus and packing-density convention and are excluded from our volumetric density/modulus evaluation. Foam entries describe effective compressive properties. The appearance of values in a simulator-facing column does not establish that a reconstructed object's scale, attachments, loading, or constitutive law have been recovered. Our downstream models predict agreement with the released material prior under the masks and applicability rules in Section 3; the illustrative scenes in Appendix I are a separate interface demonstration.

E CURATION VERIFICATION AND EVIDENCE PROVENANCE

The author-provided construction account describes an iterative procedure, not an independent physical ground-truth dataset. Its screened corpus contains 21,117 grade-level records; 58 sources are cited, 84 candidate grades are named, and 70 have a complete (ρ, E, ν) triple. These are *corpus* counts. The *exported PartNeXt annotation* uses ten classes, 23 nonempty bucket identifiers, and 26 distinct complete class/property tuples. We independently recounted the latter from the released CSV with the script and JSON report included in the artifact. This distinction matters because 70 available corpus grades do not imply 70 distinct targets in this benchmark.

Internal checks. The supplied protocol specifies eight heterogeneous suites spanning schema and enum validity, primary-key resolution, interval containment, derived-quantity consistency, grade single-valuedness, source closure for corpus-backed quantities (with an explicit reference-range exception), and an adversarial six-pattern self-test. A 457-entry table maps part names to sets of admissible material classes and covers 691 of 741 category/name combinations, representing 312,311 released parts. It is a plausibility screen, not a gold label for the admissible choices. The protocol also describes 13 simulator-facing criteria for each of 6,011 parts routed to volumetric deformation, for 78,143 checks, plus object-level coupling checks on 3,485 objects. The authors report zero failures after their rules reached a simultaneous fixed point. The admissibility table, assertion logs, integrator trajectories, and repair history were not supplied with the release we evaluated, so those zero-failure outcomes remain author-reported. Even if reproduced, self-consistency and class admissibility would not establish that each individual part has the correct material grade.

Independently recountable release properties. Our CSV-only audit found 333,152 unique model/part keys, 23,221 objects, 50 categories, 329,451 complete property rows, 3,701 abstentions, 26 distinct complete class/property tuples, and no nonmissing point outside its own recorded interval. Solver-route totals are 282,263 rigid, 41,177 cloth/shell, 5,758 MPM elastic, 253 FEM elastic, and 3,701 abstentions. There are 41,177 fabric rows carrying membrane-modulus/packing-density caveats, 5,378 near-incompressibility caveats, and 190 effective-foam caveats. This recount verifies values and relationships visible in the CSV; it does not recreate the authors' visual or load-case judgment.

Reported human review. The authors state that they randomly sampled approximately 5% of the final released parts, inspected each sampled mesh or render, and assessed the plausibility of its material class and assigned property prior. They report a 99.7% pass rate. The exact denominator, review rubric, sampled part IDs, per-part judgments, adjudication, and corrected records were not retained. We consequently present this as an author-reported process statistic, without a derived confidence interval, category breakdown, or independent test-set claim. Manual plausibility review of a sample is compatible with pipeline-authored labels: the sample was checked by people, but the released annotations were generated by the rules rather than manually assigned part by part. Appearance alone cannot verify an asset's actual modulus or density.

What the load screens changed. The supplied account describes three class-fallback failures detected using an idealized governing load. A reference-range grade replaced the old fallback for 169 upholstery-foam rows (modulus 0.00369 to 0.00004 GPa), 325 caster-wheel rows (0.000646 to 0.040

Table 6: Exploratory train-defined category–name ambiguity audit on the ID test split. A key is unique if training assigns it one coarse material class and ambiguous if training assigns it more than one; unseen keys (20 parts) are omitted. Log-modulus targets are applicable for 3,176 unique-key and 4,997 ambiguous-key parts. Metrics are object-balanced means across three fitted seeds, except the deterministic lookup. The stratification uses no test labels.

Stratum	Predictor	Objects	Parts	F1 ↑	$\log E$ ↓
Unique	Category–name lookup	435	3,410	68.6	0.086
	Semantic independent	435	3,410	68.2	0.088
	Fused independent	435	3,410	67.0	0.106
	Fused set attention	435	3,410	67.0	0.102
Ambiguous	Category–name lookup	597	5,353	58.4	0.650
	Semantic independent	597	5,353	61.1	0.615
	Fused independent	597	5,353	59.0	0.550
	Fused set attention	597	5,353	61.0	0.567

GPa), and 92 plasticized-PVC cable-jacket rows (2.25 to 0.025 GPa). The first revised idealized 200-mm cushion compression from 0.076 to 7.0 mm under 1.4 kPa; the wheel screen revised a 10-mm tread's linear strain from 155% to 2.5% under 1 MPa; the cable screen revised a 300-mm cantilever's tip deflection from 1.3 to 134 mm. These 586 rows are directly identifiable in the CSV by a reference-range interval and an empty source-ID field. The calculations expose grossly implausible fallbacks within chosen load cases. They are not measured deformations of corresponding 3D assets and do not validate the authors' simulator route on an external test set.

F Scope and dataset limitations

The released priors inherit the annotation pipeline's semantic preferences and the underlying digital assets' coverage biases. Material classes are highly imbalanced. Some small or internal parts are invisible in the selected views, and normalized geometry does not reveal absolute size, internal fill, manufacturing process, anisotropy, or a complete constitutive law. Source-linked numerical values can support plausible asset initialization, but they do not identify an individual object's measured mechanical response. A density prediction alone does not validate total mass without reliable volume and scale; an elastic-modulus assignment alone does not validate deformation without boundary conditions, geometry, and suitable material models.

Fabric illustrates this limitation sharply. The source registry contains both membrane and bulk textile grades, but the exported fabric parts use a single membrane/packing triple; even the 1,066 class-range fallback rows carry the same point values. Pillows, mattresses, blankets, and upholstery therefore cannot be distinguished by fabric properties in the present release. A fabric surface label does not reveal a pillow's fill or a mattress's layered mechanics. We retain the released labels for reproducibility and do not treat simulator contact with an authored proxy as validation of those identical fabric values.

A release-key recount against the cached upstream masks and hierarchy partitions the 16,993 upstream parts without released rows into 2,711 in wholly unannotated models, 9,844 named masks with at most two faces, 12 masks without a hierarchy name, and 4,426 named masks with more than two faces. Nine of the unnamed masks also have at most two faces; the categories are counted once in the stated order. The companion source's earlier 14,044/2,949 decomposition is inconsistent with this release-specific join and is superseded here (Appendix E). We do not fill these gaps with pseudo-ground-truth during evaluation. Source measurement, reference ranges, class imputation, and asset-specific physical measurement are distinct forms of evidence. This benchmark preserves those distinctions and does not claim calibrated physical uncertainty from matching source intervals.

G Additional downstream results

Where semantic keys are ambiguous. Table 6 separates test parts by whether their exact category and lowercased part name has one or multiple coarse material classes in training. This is an exploratory analysis selected after examining the main results; the stratum itself uses training labels only. Of

Table 7: Exploratory disjoint-name completion on train-defined category–name strata. Queries remain fixed from zero to three observed labels; columns are object-balanced percentages, averaged over three observation seeds. Set attention also averages three fitted seeds. The table baseline uses train-only counts and validation-selected strengths. Strata use no test labels; pairwise intervals are reported in text.

Key	Predictor	Objects	Acc. $b = 0$	Acc. $b = 3$	F1 $b = 0$	F1 $b = 3$
Ambiguous	Support-aware lookup	460	74.7	76.4	60.2	61.0
	Set attention	460	72.9	75.2	60.7	62.5
Unique	Support-aware lookup	340	97.1	97.2	67.9	68.0
	Set attention	340	95.0	94.9	65.5	65.4

Table 8: Post-hoc test-time descriptor permutation on all 634 ID-test objects. Each part's visual and/or geometric descriptors are rotated among test parts with the same category and lowercased name; 8,716 of 8,783 parts move, and 8,193 parts have applicable log-modulus targets. Semantic, whole-object visual, structural, and target fields and fitted weights remain fixed. Differences are permuted minus original; positive log-modulus error means degraded numerical prediction. Intervals are paired 95% geometry-group bootstrap intervals over 2,000 resamples, averaged across three fitted seeds.

Predictor	Swapped fields	Δ F1 (points)	$\Delta \log E$ (dex)
Independent	Part visual	$+1.22$ $[-0.63, 3.04]$	$+0.033$ $[0.013, 0.052]$
Independent	Part geometry	-1.40 $[-2.89, 0.11]$	$+0.033$ $[0.018, 0.049]$
Independent	Both	-0.44 $[-2.47, 1.51]$	$+0.056$ $[0.033, 0.079]$
Set attention	Part visual	$+0.33$ $[-1.49, 2.25]$	$+0.025$ $[0.007, 0.043]$
Set attention	Part geometry	-1.72 $[-2.89, -0.62]$	$+0.031$ $[0.019, 0.044]$
Set attention	Both	-1.19 $[-3.84, 1.23]$	$+0.050$ $[0.027, 0.073]$

8,783 ID-test parts, 3,410 have unique keys, 5,353 have ambiguous keys, and 20 have unseen keys. On ambiguous keys, fused independent prediction reduces object-balanced log-modulus error from 0.615 to 0.550 dex relative to the semantic-only independent model. The paired fused-minus-semantic difference is -0.065 dex (geometry-group bootstrap 95% interval $[-0.130, -0.004]$). It does not improve material F1: the difference is -2.07 points, with interval $[-4.67, 0.41]$. On unique keys, category–name lookup remains strongest and fusion lowers F1. The analysis locates a possible numerical use for visual/geometric evidence where names are ambiguous, but it does not establish visual material recognition, an architecture advantage, or independent mechanical accuracy. The saved prediction audit script reproduces all strata, metrics, and intervals.

Conditional completion on ambiguous keys. On the train-defined ambiguous-key subset, three observed labels raise set-attention F1 by 1.84 points (paired 95% interval $[0.46, 3.24]$) and the support-aware table by 0.78 points ($[-1.24, 2.69]$). The difference between these gains is 1.06 points ($[-1.24, 3.30]$), so a larger learned contextual gain is not established. At three labels, set attention is 1.49 F1 points above the table ($[-1.52, 4.65]$), also unresolved. Unique-key parts strongly favor the table. The subgroup is post-hoc and its target labels are pipeline-authored; these numbers do not validate material truth or a general learned-context advantage.

Does the fitted model use part-specific evidence? To separate a model's use of part descriptors from the correlations in its semantic keys, we deterministically rotate part-level visual and geometry descriptors within each exact category–name group of the ID test set. This intervention leaves test targets, names, semantic embeddings, global object appearance, hierarchy, and fitted model weights unchanged. Singleton keys are left intact. On the complete test population, visual-only and geometry-only permutations each raise the fused independent model's log-modulus error by 0.033 dex; swapping both raises it by 0.056 dex. Table 8 reports all paired differences and intervals for the independent and set-attention predictors. Geometry perturbation also lowers set-attention material F1 by 1.72 points, while the other F1 changes have intervals crossing zero. Thus the fitted predictors use some part-specific visual and shape evidence for agreement with released numerical priors, even though fusion fails to beat semantic-only classification in the main benchmark. This is a post-hoc

Table 9: Nested training-object fractions for the fused independent predictor, with 4,400 allocated updates per fit and validation-selected checkpoints (three seeds). Per-category rounding yields the stated actual object counts. Numerical errors use the same reference-valid test queries for every fraction.

Fraction	Objects	F1 (%)	Log-ρ MAE	Log-E MAE	ν MAE
10%	293	57.9±0.9	0.206±0.005	0.491±0.003	0.022±0.000
25%	722	62.3±0.5	0.175±0.003	0.427±0.007	0.019±0.000
50%	1411	64.0±1.7	0.167±0.004	0.381±0.007	0.017±0.000
100%	2802	65.5±1.7	0.157±0.004	0.373±0.011	0.016±0.001

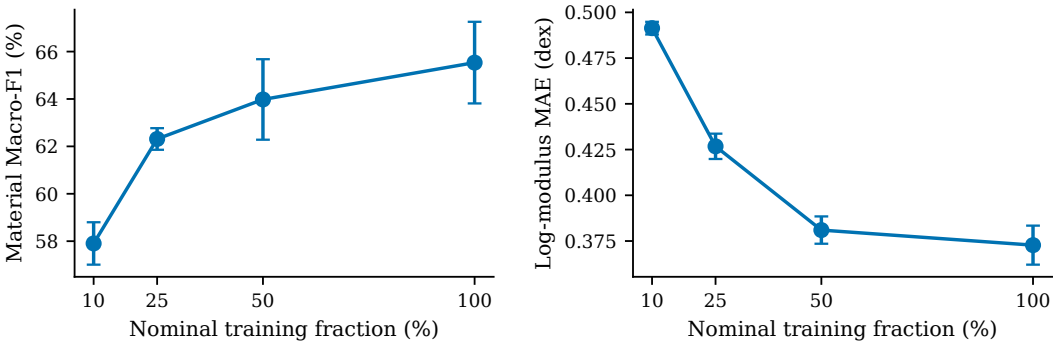

Figure 3: Annotation quantity with equal allocated optimization budgets. Nested subsets contain 293, 722, 1,411, and 2,802 training objects. Each fit receives 4,400 updates; validation selects its checkpoint. Error bars show sample SD across three seeds of the fused independent predictor.

input-dependence diagnostic, not evidence that the inferred priors match measured properties. CPU replay of saved GPU checkpoints changes at most three of 8,783 prototype predictions per fit; the paired perturbation compares original and swapped descriptors under the same CPU backend.

Category holdout. The five-category holdout remains challenging (Table 10). Name-based lookup attains 72.5% accuracy and 46.0% ten-class F1, exceeding the learned fused models on both metrics. Set context improves accuracy over the independent model, but not Macro-F1. Transferable name associations therefore remain a strong reference even on unseen categories; these rankings alone do not identify why the learned models degrade. Two of the ten material classes have no reference support and contribute zero to ten-class Macro-F1. Over the eight supported classes, lookup F1 is 57.6%, versus 55.6%, 53.8%, and 54.0% for fused independent, set, and relation models (means over three fitted seeds), so the ranking does not depend on the absent-class convention. Absolute ten-class F1 levels should not be directly compared with the in-distribution split.

Numerical targets and prediction heads. Classification followed by lookup is a strong parameterization for this release's finite property vocabulary (Table 11). For the independent predictor, lookup achieves log-density/log-modulus errors of 0.156/0.364 dex, compared with 0.215/0.518 for the joint regression head. The latter predicts nonphysical Poisson ratios for approximately 0.6% of scored parts. This is evidence about the declared joint objective and classification-selected checkpoints; it does not establish that lookup is superior to every numerically tuned regression system or for unseen physical materials.

Selecting the next annotation. On the 100-object active subset, final F1 is 67.5% with random acquisition, 67.9% with entropy, 67.9% with feature diversity, and 67.6% with expected-entropy look-ahead. The additional hypothetical predictions of look-ahead do not yield a consistent advantage. This experiment supports using inexpensive acquisition rules as references; it does not support a new active-learning algorithm claim.

Table 10: Category holdout: Chair, Knife, Mug, Laptop Computer, and Skateboard. Models train and select checkpoints using only the other categories. Lookup keys back off using training statistics. Material metrics use all ten classes, including zero for absent classes. Learned values are mean ± SD over three seeds.

Predictor	Accuracy (%)	F1 (%)	Log-E MAE
Name lookup	72.5	46.0	0.717
Category–name lookup	72.5	46.0	0.717
Independent	63.5±3.7	44.5±1.7	0.731±0.018
Set attention	66.2±0.9	43.1±0.9	0.695±0.032
Relation bias	65.9±1.8	43.2±2.1	0.726±0.018

Table 11: Numerical parameterization under shared inputs and classification supervision (three seeds). Regression adds masked, training-scale-normalized Smooth-L_1. Invalid templates receive a training-median fallback without removing reference-valid queries. Nonphysical ν is outside $(-1, 0.5)$ or nonfinite, counted over all 8,783 queried parts. Numerical columns are MAE in the indicated coordinate; SDs round to three decimals.

Predictor	Head	Log-ρ MAE	Log-E MAE	ν MAE	Invalid ν (%)
Independent	Lookup	0.156±0.001	0.364±0.006	0.016±0.000	0.0±0.0
Independent	Regression	0.215±0.014	0.518±0.037	0.025±0.002	0.6±0.5
Relation bias	Lookup	0.163±0.003	0.377±0.008	0.017±0.001	0.0±0.0
Relation bias	Regression	0.238±0.008	0.553±0.029	0.027±0.002	0.4±0.1

Input controls and parameter counts. The tables expose both descriptor choices and active trainable model size. Modality changes alter the available evidence and the input encoder size; only comparisons within the same descriptor regime approximately match trainable parameters.

Table 15 gives the reference support behind the ten-class average. In the ID test, the numerical metrics retain 612 objects with at least one applicable part for each quantity: 8,193 parts for density/modulus and 8,579 for Poisson ratio. Material classification retains all 634 objects and 8,783 parts. Table 16 separately records invalid-template fallback frequencies. These queries remain in numerical error evaluation; choosing an inapplicable template cannot improve a method by excluding its error.

Matched pretrained-transfer uncertainty. Both pretrained runs complete their entire requested populations without a failed object. The Qwen comparison resamples 600 geometry groups covering all 634 test objects; the VoMP comparison resamples 99 groups covering its 100 selected objects. Paired intervals use the same 2,000-resample specification as the core experiments. Relative to Qwen, the visual-only and semantic-only F1 gains are 15.88 points ($[11.99, 19.76]$) and 28.97 points ($[25.31, 32.80]$). Relative to the VoMP adapter, visual-only learning reduces log-density error by 0.244 dex ($[0.175, 0.309]$), log-modulus error by 0.518 dex ($[0.377, 0.647]$), and Poisson-ratio error by 0.0138 ($[0.0085, 0.0191]$). These are pointwise comparisons on released priors, conditional on the fitted models and declared pretrained interfaces; rendering, available evidence, and supervision are not held equal. Complete paired comparisons for all five learned controls are included in the result archive.

Exact semantic-key coverage. This post-hoc train-only diagnostic uses the lowercased pair (object category, part name) as an exact key. Among 8,783 ID-test parts, 8,763 (99.77%) have a key present in training; 3,410 (38.83%) have a key with only one coarse class in training, and 2,798 (31.86%) have only one fine property prototype. The released test class occurs among its training key's classes for 8,695 parts (99.00%), and its fine prototype occurs for 8,579 (97.68%). These are coverage and ambiguity diagnostics, not accuracy against independent material labels. The object-balanced fractions and a machine-readable report are included with the analysis artifact.

Exploratory exact-key supervision holdout. After reviewer feedback, we froze a second protocol before model fitting. For each exact key formed by object category, a separator, and the lowercased, stripped original part name, the first 64 bits of `SHA256("key-holdout-v1-s0:" + key)` assign it to train $[0, .70)$, validation $[.70, .85)$, or test $[.85, 1)$. The original geometry-group object

Table 12: Choosing observations on 100 fixed eligible test objects, each with four to eight candidates and disjoint query-name groups. Policies reuse the same relation-biased completion model and fixed queries. F1 is mean ± SD over three fitted models. Time is approximate wall-clock seconds per object for a complete three-query trajectory on the shared GPU server, including prediction; it is not an isolated hardware benchmark.

Query policy	F1 at $b = 1$	F1 at $b = 3$	Seconds/object
Random	67.6±2.0	67.5±1.3	0.010
Posterior entropy	67.8±2.2	67.9±1.9	0.007
Feature diversity	67.8±2.2	67.9±2.1	0.007
Query look-ahead	67.4±1.9	67.6±1.6	0.020

Table 13: Modality and explicit-material-word controls for the relation-biased predictor (three seeds). T1 and T2 use their separately defined populations. Scrubbing removes a fixed material-word list from original semantics before feature hashing; it does not guarantee removal of all semantic clues. Material metrics are object-balanced percentages.

Task	Descriptors	Accuracy ↑	F1 ↑
T1	Fused semantics	81.5±0.9	65.3±0.6
T1	Material-word scrubbed	81.8±0.2	65.0±0.8
T2, disjoint names, $b = 3$	Semantics	83.4±0.2	68.0±0.5
T2, disjoint names, $b = 3$	Visual/geometric	75.0±0.9	56.5±0.7
T2, disjoint names, $b = 3$	Fused	82.5±0.6	65.6±0.5

split is preserved. Within each object, only parts whose key is assigned to that object's split enter the supervised loss or metric; other parts are visible, unlabeled context. The resulting supervised sets have 456/80/79 distinct keys and 29,679/1,007/1,557 queried parts in 2,719/347/296 objects. Assertions confirm zero supervised-key overlap and matched baseline/model test queries. Of the 79 test keys, 77 appear in unlabeled training-object context, and the class/prototype vocabulary remains public; thus the test withholds exact-key *supervision*, not semantic tokens. The train-only lookup counts prototypes with equal total weight per training object and backs off from name to category to global; its exact-key branch cannot fire on test. We additionally report a global and a category-only control. Three fit seeds use the same grouped object splits. Paired geometry-group bootstrap uses 1,500 draws and conditions on the fixed 79 keys and fitted seeds. A separate 1,500-draw paired key-cluster bootstrap resamples those keys with replacement, carries all queried parts for each key together, and retains equal total weight for every active object in a draw. The latter is a sensitivity analysis over test-key composition, not a new hash-partition study. Ten-class F1 is unstable for rare classes: validation lacks composite, foam, and leather, while test has only 1/0/2 parts in these classes (Table 19). These results evaluate agreement with authored assignments, not independent material correctness. The exact dataset hash, masks, predictions, and paired analyses accompany the artifact.

For the train-only backoff, 262 of 1,557 test queries use the name table and 1,295 use the category table; none reach the exact-key or global branch. A capacity control changes only the semantic MLP width to 232 (499,755 parameters), close to the fused MLP's 502,267; the default semantic MLP has 300,283. The capacity control retains the same optimizer, validation rule, grouped split, and three fit seeds. This comparison controls model size approximately, while leaving possible modality-specific inductive biases. Two additional salts were fixed in an analysis plan before their model fits; Table 20 reports all three partitions without selection. The fused F1 point gain over the size-matched semantic model is positive in each partition, but its magnitude varies from 4.46 to 12.53 points, and the $s2$ accuracy gain is only 1.31 points. Foam has no test support under any salt; $s1$ also lacks composite and leather. Three exploratory partitions are insufficient for a population-level split interval, and none provides independent label truth.

G.1 Additional annotations improve the fitted prior

At equal update budgets, more annotated objects improve the fitted material prior (Appendix Figure 3). Mean F1 rises from 57.9% at 10% training data to 62.3%, 64.0%, and 65.5% at 25%, 50%, and 100%. Full-versus-10% F1 improves by 7.63 points (95% interval $[4.96, 10.47]$); log-modulus error falls

Table 14: Active trainable parameters for the fused unconditional comparison. The frozen DINOv2 backbone is shared and excluded. Widths approximately match a 128-wide set predictor; relation bias adds twelve parameters in each of two blocks.

Predictor	Width	Parameters	Input dimensions
Independent	160	502,267	1817
Mean pooling	152	490,683	1817
Fixed graph	144	517,707	1817
Set attention	128	501,659	1817
Relation bias	128	501,683	1817

Table 15: ID-test reference support and class-specific F1 from the object-normalized confusion matrix (mean $\pm$ SD over three fitted seeds). Supports count parts, whereas the metrics weight objects equally. Very small class supports limit interpretation of individual class estimates.

Material	Test parts	Semantic independent	Visual independent	Fused set
ceramic	217	85.4±0.4	82.3±1.6	87.0±1.0
composite	17	68.1±0.0	53.6±8.8	64.7±3.7
fabric	386	91.7±1.7	78.4±1.0	92.7±0.9
foam	6	0.0±0.0	0.0±0.0	0.0±0.0
glass	159	74.0±1.3	43.8±1.7	71.7±2.6
leather	9	28.9±8.3	4.3±7.4	30.6±4.7
metal	2972	84.2±0.2	78.6±0.6	83.9±0.6
plastic	3728	77.4±1.0	70.5±0.9	77.8±1.1
rubber	199	86.5±1.5	64.3±0.5	84.7±1.2
wood	1090	81.1±2.4	70.6±0.7	77.6±1.3

from 0.491 to 0.373 dex, a reduction of 0.119 ($[0.085, 0.154]$). Full-versus-50% intervals include zero. This is one frozen-feature learner and nested subset selection, not a universal scaling law. Table 9 reports all quantities.

Pretrained transfer. Zero-shot Qwen reaches 38.8% F1 on all 634 ID-test objects, below visual-only learning (54.6%) and semantic-only learning (67.7%). On 100 preselected objects, adapted VoMP also has higher numerical-prior errors than in-domain fusion and category–name lookup (Appendix Tables 17–18). Interfaces differ; these comparisons do not isolate architectures or validate measured mechanics.

Table 16: Invalid-prototype fallback frequency in unconditional prediction (percent, mean ± SD over three fits). Each entry is the fraction of reference-valid query parts whose predicted prototype lacks that applicable quantity, receiving a training-median value. Denominators are 8,193 for density/modulus and 8,579 for Poisson ratio. These are part frequencies, not object-balanced errors. S and V follow Table 2.

Predictor	Input	ρ fallback	E fallback	ν fallback
Independent	S	3.87±0.11	3.87±0.11	0.35±0.11
Independent	V	2.03±0.88	2.03±0.88	0.51±0.30
Independent	V+S	0.81±0.11	0.81±0.11	0.28±0.02
Mean pooling	V+S	0.80±0.08	0.80±0.08	0.26±0.04
Fixed graph	V+S	0.78±0.04	0.78±0.04	0.24±0.01
Set attention	V+S	0.76±0.12	0.76±0.12	0.28±0.03
Relation bias	V+S	1.45±0.86	1.45±0.86	0.29±0.02

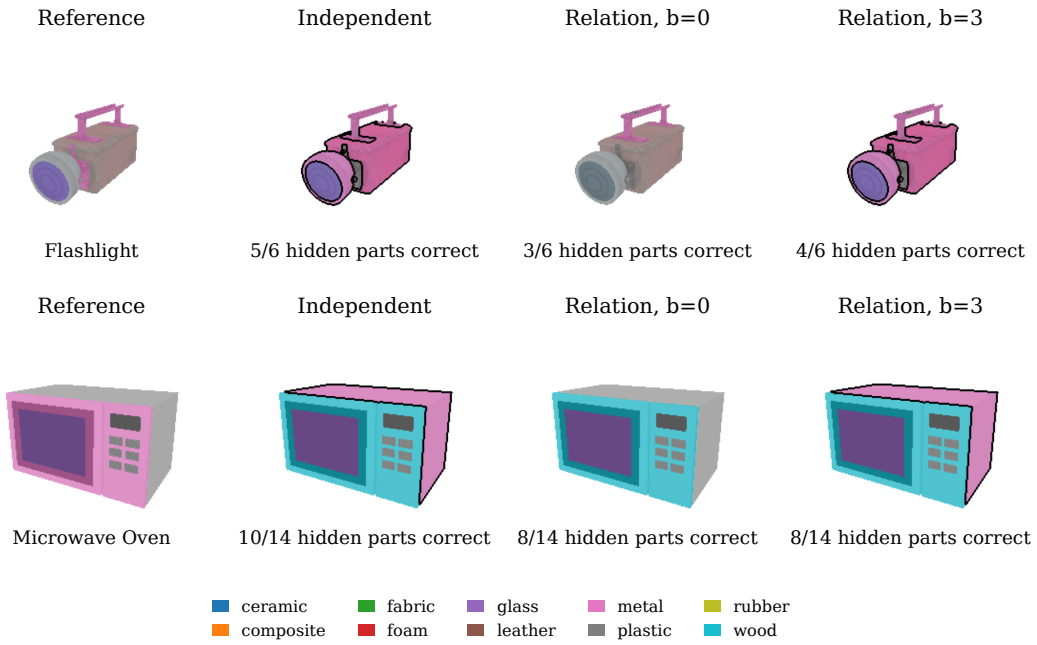

Figure 4: Two examples selected by an independent object hash before inspecting predictions, using disjoint-name episodes and training/observation seed zero. Query parts are colored by material; observed supports have black boundaries; unscored context is gray. Counts score all hidden parts, including those invisible in this view. The relation model improves by one hidden part on the flashlight and by none on the microwave; independent prediction remains stronger on both examples. These illustrations are not a performance-selected success gallery.

Table 17: Material classification on the same 634 successfully processed held-out objects. Qwen receives an outlined view, a target crop, and category/name/path text, with no in-domain updates; it ranks ten single-letter next-token choices. Visual-input learned methods use frozen six-view descriptors; semantic controls use category/name/path features, and all learned methods receive in-domain supervision. This is a pretrained transfer comparison with different backbones and inputs. Metrics are object-balanced percentages.

Method	Accuracy ↑	F1 ↑
Qwen2.5-Omni-7B (zero shot)	56.1	38.8
Category–name lookup	82.9	65.0
Semantic independent	82.5±0.3	67.7±1.1
Visual independent	74.4±0.4	54.6±1.6
Fused independent	81.8±0.1	65.4±1.1
Fused set attention	82.1±0.5	67.1±0.6
Fused relation bias	81.5±0.9	65.3±0.6

Table 18: Agreement with numerical priors on the same 100 successful objects from a preselected 100-object test subset; 96 objects have applicable targets for each quantity. Selected controls are shown; the artifact contains all evaluated rows. VoMP uses its released weights and 150 adapted views; the other models use in-domain labels. The renderer and Torch backend adaptations are documented in the appendix. This comparison does not measure asset-specific mechanics. Log errors are in dex and Poisson-ratio error is dimensionless; lower is better.

Method	Log-ρ MAE	Log-E MAE	ν MAE
VoMP (rendering/backend adapter)	0.481	1.081	0.040
Category–name lookup	0.159	0.303	0.015
Semantic independent (lookup)	0.184±0.031	0.346±0.029	0.017±0.001
Visual independent (lookup)	0.237±0.010	0.563±0.037	0.026±0.002
Fused independent (lookup)	0.177±0.016	0.315±0.013	0.016±0.001
Fused relation bias (lookup)	0.180±0.013	0.322±0.011	0.017±0.001
Fused independent (regression)	0.249±0.011	0.470±0.024	0.024±0.001
Fused relation bias (regression)	0.292±0.020	0.541±0.054	0.027±0.003

H Related task interfaces

Table 21: Selected part-material task interfaces. The targets and supervision protocols differ; the comparison is descriptive.

Work	Representation	Prediction or query
PartNeXt (Wang et al., 2025)	Triangle parts; semantic tree	Fine-grained structured part annotation
3DCoMPaT200 (Ahmed et al., 2024)	Styled 3D parts; part-material labels	Compositional recognition and retrieval
HG-CAD (Bian et al., 2023)	CAD face, body, and assembly graphs	Body material recommendation given labels on other bodies
VoMP (Dagli et al., 2026)	Visual geometry; volumetric queries	Density, modulus, and Poisson-ratio fields
Gen2Physics (Comi et al., 2026)	Mesh material components and rendered views	Material masks plus component density and structure
PartNeXt-Mat	PartNeXt triangle parts and tree; material priors	Hidden-part prediction with 0, 1, or 3 observed labels; fixed queries and name-disjoint control

These works do not share an identical evaluation target. Our fixed-graph baseline does not reproduce HG-CAD; Gen2Physics and 3DCoMPaT200 are related resources, not numerical baselines in our suite.

Exploratory cross-resource transfer. As a diagnostic beyond the PartNeXt-Mat release, we fixed a hash-ranked sample of 50 public UniPhys-Bench objects before reading their annotations and

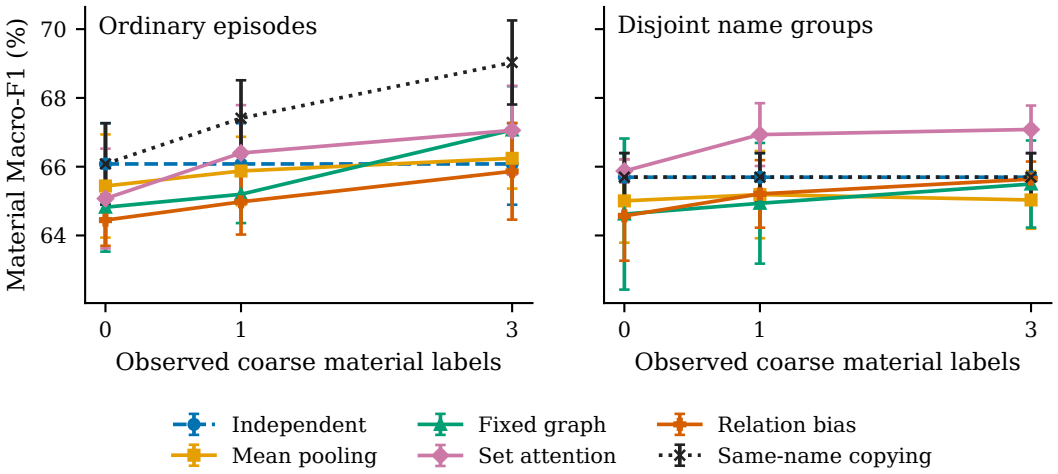

Figure 5: Fixed-query completion (supplementary visualization). Each curve scores the same hidden parts; bars show SD over three fits. The ordinary and disjoint-name populations differ. Semantic and support-aware controls are in Table 4.

Table 19: Support of the exploratory exact-key-supervision holdout and object-balanced class F1 for the fused independent model, averaged across three fits. Zero-support test classes contribute zero to the ten-class mean; the supported-nine-class mean is 34.5%.

Material	Val parts	Test parts	Test F1 (%)
Ceramic	67	7	11.4
Composite	0	1	0.0
Fabric	40	49	50.0
Foam	0	0	0.0
Glass	96	2	18.6
Leather	0	2	0.0
Metal	439	546	75.6
Plastic	193	733	50.0
Rubber	29	40	51.6
Wood	143	177	53.2

retrieved only annotation JSON through hf-mirror. Forty-five objects had complete referenced part files; five returned missing-file responses. We take the prefix of each external material string and map `ceramics`→ceramic and `cloth`/`textile`→fabric; other strings must match a class exactly. Five out-of-vocabulary parts are excluded. The 375 mappable parts cover six classes: ceramic 3, glass 9, metal 132, plastic 54, rubber 8, and wood 169. With zero target-domain updates or model selection, three frozen semantic-only PartNeXt-Mat fits average 31.0% object-balanced accuracy and 10.2% ten-class F1 (17.0% over supported classes); train-only name/category/global backoff reaches 31.8%, 5.0%, and 8.3%, respectively. A paired 2,000-draw object bootstrap gives a semantic-minus-backoff F1 difference of 5.21 points, 95% interval $[3.33, 7.10]$, while the accuracy difference is -0.84 points, $[-7.13, 5.61]$. This remains weak transfer, not external validation of PartNeXt-Mat labels or physical properties: category/name conventions differ, the hierarchy path is unavailable, and the external labels are reference annotations rather than measured mechanics. The pinned selection, missing IDs, per-object confusion matrices, and seeds accompany the analysis artifact (Li et al., 2026).

I Illustrative image-to-mesh scene authoring

Figure 6 illustrates scene authoring with an image-derived asset. We use a real Anthurium photograph and its foreground mask from the CC BY 4.0 Botany-Bot dataset (Adebola et al., 2025), downloaded as two individual files through hf-mirror. We reconstruct a colored 3D mesh with the pretrained SAM 3D Objects model (SAM 3D Team et al., 2025). A deterministic height-and-color rule partitions

Table 20: All three fixed exact-key-supervision hash partitions. Object-balanced accuracy and ten-class F1 are percentages; fitted rows show mean ± SD across three seeds. This is descriptive split sensitivity, not an independent material-label test.

Salt	Keys/parts	Predictor	Accuracy ↑	F1 ↑
s0	79/1557	Global prior	57.8	7.3
		Category prior	46.0	15.4
		Name/category backoff	54.5	21.6
		Size-matched semantic	57.8±1.7	23.4±0.8
		Fused independent	63.9±0.8	31.0±1.9
s1	77/730	Global prior	48.2	6.5
		Category prior	66.6	28.9
		Name/category backoff	57.6	18.6
		Size-matched semantic	60.7±6.1	27.7±6.4
		Fused independent	68.1±1.2	40.3±1.1
s2	80/1073	Global prior	58.1	7.3
		Category prior	58.7	22.4
		Name/category backoff	59.5	29.2
		Size-matched semantic	66.4±1.6	37.1±3.1
		Fused independent	67.7±3.1	41.5±2.8

Table 22: Component roles for Figure 6. The photographed plant has no PartNeXt-Mat part annotations; visual components therefore have no claimed (ρ, E, ν) prediction. The pot's numerical triple is an illustrative manual match to a released thermoplastic grade, not a measurement or a solver calibration.

Component	Representation and assigned value
Pot	Rigid collider; illustrative manual reference $\rho = 1200\,\mathrm{kg/m^3}$, $E = 2.25\,\mathrm{GPa}$, $\nu = 0.364$.
Stems	Three hinge capsules, each $0.015\,\mathrm{kg}$; $k_\theta = 0.15\,\mathrm{N\,m/rad}$ and $c_\theta = 0.025\,\mathrm{N\,m\,s/rad}$.
Soil, leaves, flower	Visual mesh regions skinned to saved hinge states; no independent deformable solver bodies or material prediction.

faces of that actual mesh into pot, soil, foliage, and flower regions. This illustrative split is not a PartNeXt annotation or a prediction of our benchmark model. The pot remains rigid; three compliant MuJoCo hinge bodies act as stem proxies, and their simulated angles drive a smooth deformation of the reconstructed plant mesh. Panels (d–f) use a stronger, explicitly authored wind episode and show rest, peak bend, and recoil at fixed camera scale. The photograph, mask, generated GLB, split GLBs, scene, 151 saved states, and rendering scripts are retained with the artifact.

The example demonstrates how a rigid container and compliant proxy stems can be assembled with a photograph-derived mesh. The geometry comes from an external reconstruction model, and the split and soft-body parameters are authored for this illustration. Crucially, MuJoCo solves only three hinge degrees of freedom and simple collision geoms: individual leaves, petals, and the high-resolution plant mesh are not deformable bodies in the solver. The displayed mesh is deterministically skinned from the saved hinge angles. The supplement includes the exact scene XML, saved joint coordinates, ball/foliage contact flags, and a time-series plot for the strong Anthurium impact. This is not a quantitative test of our material predictor, a calibrated plant mechanics simulation, or an Isaac Sim/Genesis agreement experiment. An additional AI-generated image-to-mesh plant and previous procedural flag and plant scenes remain available separately in the artifact as simulator interface prototypes.

Perturbation gallery. We additionally reconstruct four other masked real Botany-Bot plant photographs (Alocasia, Belize, Croton, and Princess Pink) with the same frozen SAM 3D model. For each of the five assets, three authored severity settings are simulated for each of three perturbations: a lateral wind gust, a sphere striking flexible stem or upper-foliage proxies, and a freely falling pot contacting the ground. This yields 45 episodes and 6,795 saved proxy states. All fifteen sphere

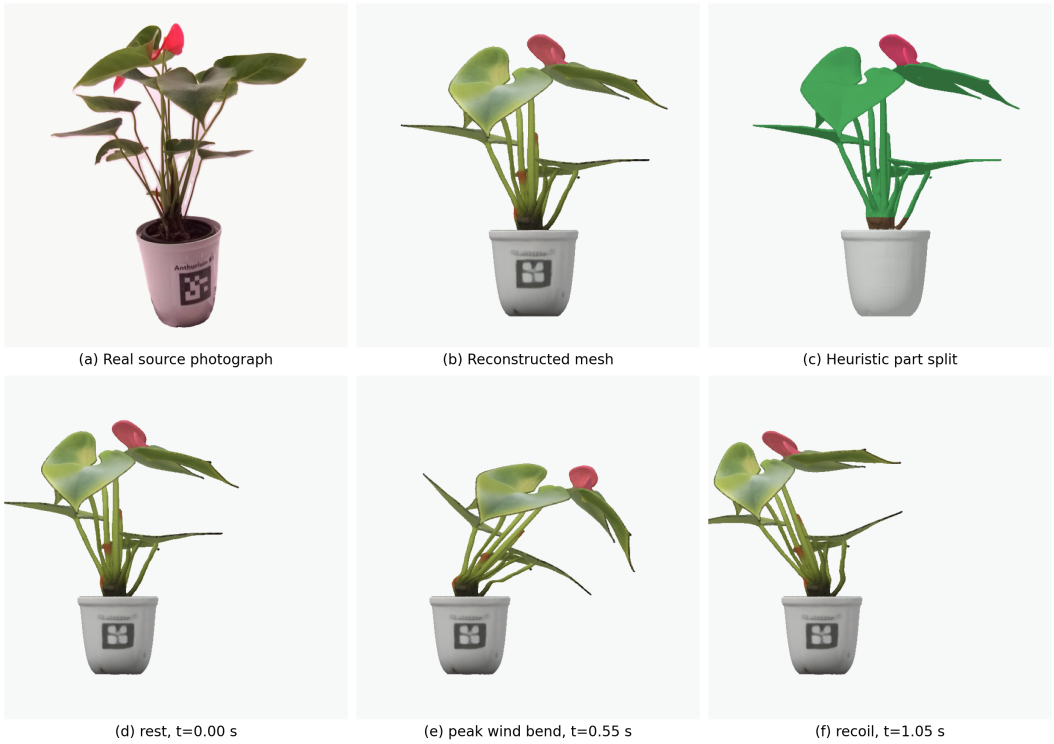

Figure 6: Real-photo-derived mixed-material scene. (a) A Botany-Bot Anthurium photograph with its supplied mask; (b) SAM 3D Objects reconstruction; (c) heuristic split of that mesh; (d–f) MuJoCo-driven visual deformation under a stronger authored wind gust. Table 22 lists the actual solver roles and authored parameters. The images are qualitative, not benchmark results.

episodes register flexible-body contact; each of the five strong sphere episodes specifically contacts an upper-foliage collider. All fifteen drop episodes register pot–ground contact. Figures 7 and 8 show a compact overview, while Figures 9 and 10 retain the strong-setting summary and one detailed sequence. All 45 videos, the other meshes, scene XML files, state arrays, contact counts, and parameters remain in the offline supplement. These episodes share an authored collider topology and do not measure how accurately material priors predict real dynamics.

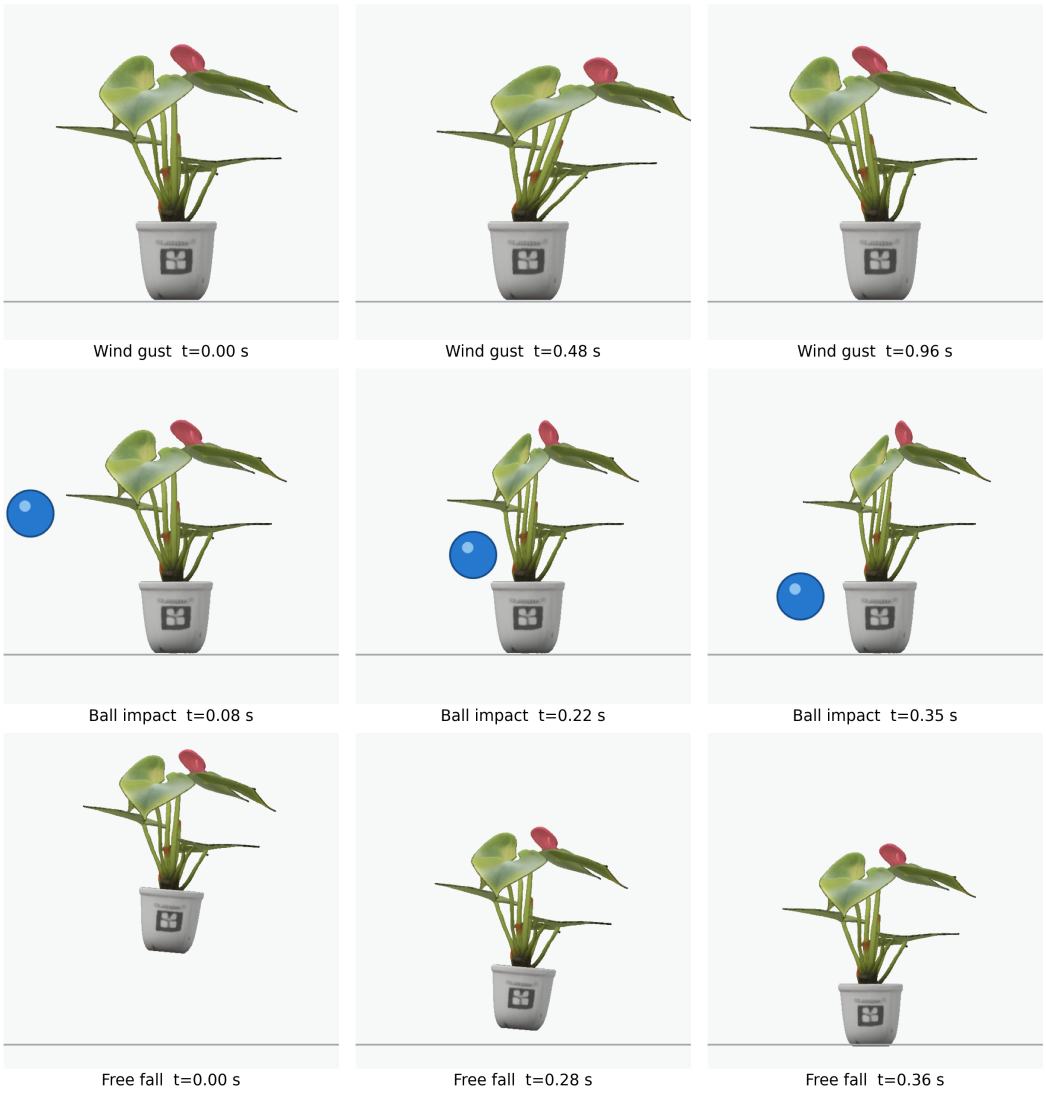

Figure 7: Three perturbations on the same photograph-derived Anthurium mesh. Rows show an authored wind gust, a simulated blue sphere contacting a stem capsule, and a rigid pot falling onto a ground plane. Columns show saved MuJoCo states at the indicated times; each scene is rendered with the reconstructed visual mesh skinned to the proxy bodies. The ball, stem capsules, pot cylinder, and floor are the actual MuJoCo collision geometry; the high-resolution mesh is visual.

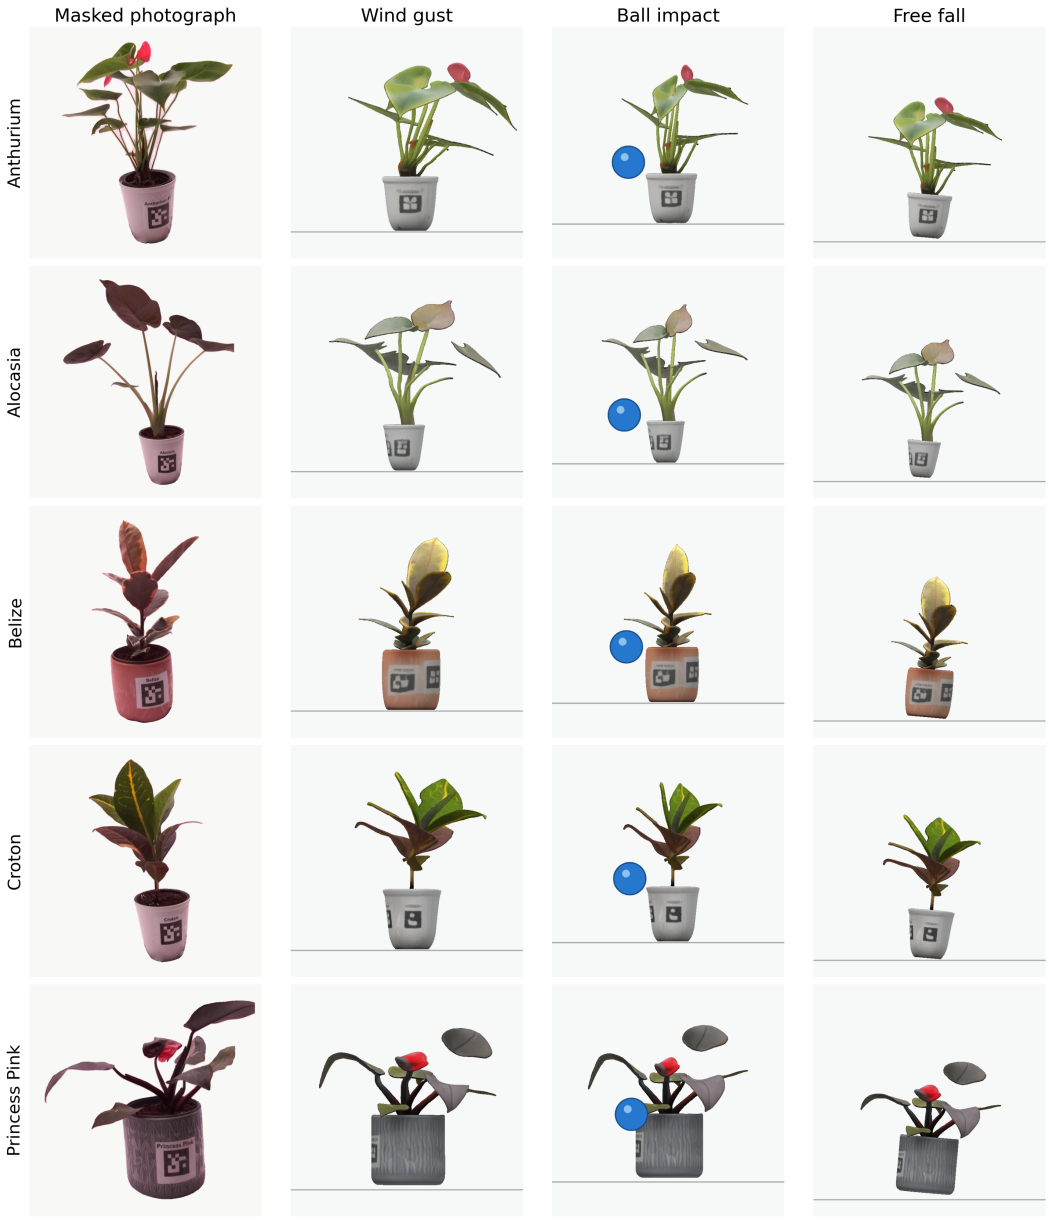

Figure 8: Selected states across five independently reconstructed real-photo plant assets. Each row begins with its masked Botany-Bot photograph, followed by one saved wind, sphere-impact, and free-fall state from its first severity setting. Two stronger settings for each perturbation are included in the artifact, giving 45 total episodes. Pot/plant face subsets and a part-property-status manifest accompany every mesh. All plant/soil (ρ, E, ν) fields remain unavailable in the released material table; visual and proxy parameters are explicitly authored. The gallery demonstrates scene creation, not measured mechanics or learned property transfer.

J VIDEO-SUPPLEMENT EVENT SEQUENCES

The HTML supplement retains all fifteen strongest-setting videos and thirty lighter episodes. Their video-frame to solver-state mappings are recorded in the reproducibility metadata. We devote the remaining appendix space to the original-GLB furniture contacts, for which native volumetric-flex states, normal contact loads, and fuller contact-to-recovery sequences can be inspected. The plant clips illustrate authored hinge-proxy scene creation, not finite-element or measured botanical response.

The accompanying offline HTML supplement contains playable videos for all 45 saved MuJoCo plant-proxy episodes, eighteen rigid-sphere/flex furniture impacts (six soft-part targets on four held-out meshes, three authored load settings each), eight single rigid-block/flex furniture impacts (four soft-part targets, two authored load settings each), one three-block/flex bed impact, and one synchronized three-run density-control video. It also contains one native MuJoCo flag-cloth/sphere impact, two earlier procedural MuJoCo previews, and a replay of a separately saved 25-state Genesis PBD flag trajectory. The site indexes the 88 separately supplied simulation-ready FEM bundles by source render, structural-anchor mask, and Genesis rest-state render. The bundle ZIP does not contain its original dynamic trajectories, so those static images are not described as FEM videos. Fifteen of the plant episodes are sphere impacts, making 43 original-GLB/plant/flag collision episodes plus three additional native-proxy control runs; the plant's foliage continues to use the three-hinge proxy described above.

K ORIGINAL-MESH HELD-OUT CASE STUDIES

The following thirteen cases use original PartNeXt test-set GLBs from the fixed 88-asset selection described in Section P. We selected recognizable assets to examine the rendered geometry, without filtering by prediction correctness. The top left panel of each figure shows the source mesh with its stored texture or vertex colors; the top right renders the same triangles and camera view with a distinct colour for each part ID. Table swatches map colours to IDs, names, predicted grades, and property values, including parts hidden in this camera view. The washing-machine GLB contains grayscale vertex colours but no embedded texture image or UV coordinates; we display its original vertex colours rather than inventing a texture. Fabric rows show the released packing-density and membrane-modulus template values with a dagger marker; these quantities are masked from volumetric evaluation and are not inputs to the flex solver. No primitive stand-in or reconstructed proxy replaces the test mesh in the upper image panels. Bed, chair, and sofa cases additionally show a rigid sphere striking the soft region. In each, MuJoCo solves contact with a pinned tetrahedral flex proxy, and the original GLB's annotated soft triangles are visually skinned from the saved proxy displacements; the rendered source triangles themselves are not the solver's collision elements. The bed, chair, and sofa baseline episodes log 509, 509, and 508 solver steps with sphere–flex contact, respectively; the nearest center surface vertex moves 0.083, 0.067, and 0.065 normalized length units downward at its peak under authored loads. The sofa target is original part ID 4, the seat cushion; the code maps that ID to the renderer's separate row index before visual skinning. The unknown physical scale of the GLB precludes a centimeter claim. These results establish solver-based proxy deformation, not physical calibration or a gain from our predictor: no predicted (ρ, E, ν) is used to set flex parameters. XML, saved vertices, part-coloured and source-appearance videos, native MuJoCo peak-state renders, and rest/peak cross-sections are included in the offline supplement. The additional light and heavy impacts change only authored sphere mass and drop height; their sofa peak center displacements are 0.044 and 0.093 normalized units, compared with 0.065 for the baseline. The separate flag example solves native two-dimensional cloth-flex vertices under a sphere impact with zero gravity to isolate contact; it is authored geometry, not a PartNeXt test asset. The remaining images assess visual plausibility of part-level predictions, not dynamic simulator accuracy. The separate 88 structural-anchor OBJ bundles have a much coarser source geometry and are documented by their original renders in the offline supplement instead of being presented as high-fidelity examples here.

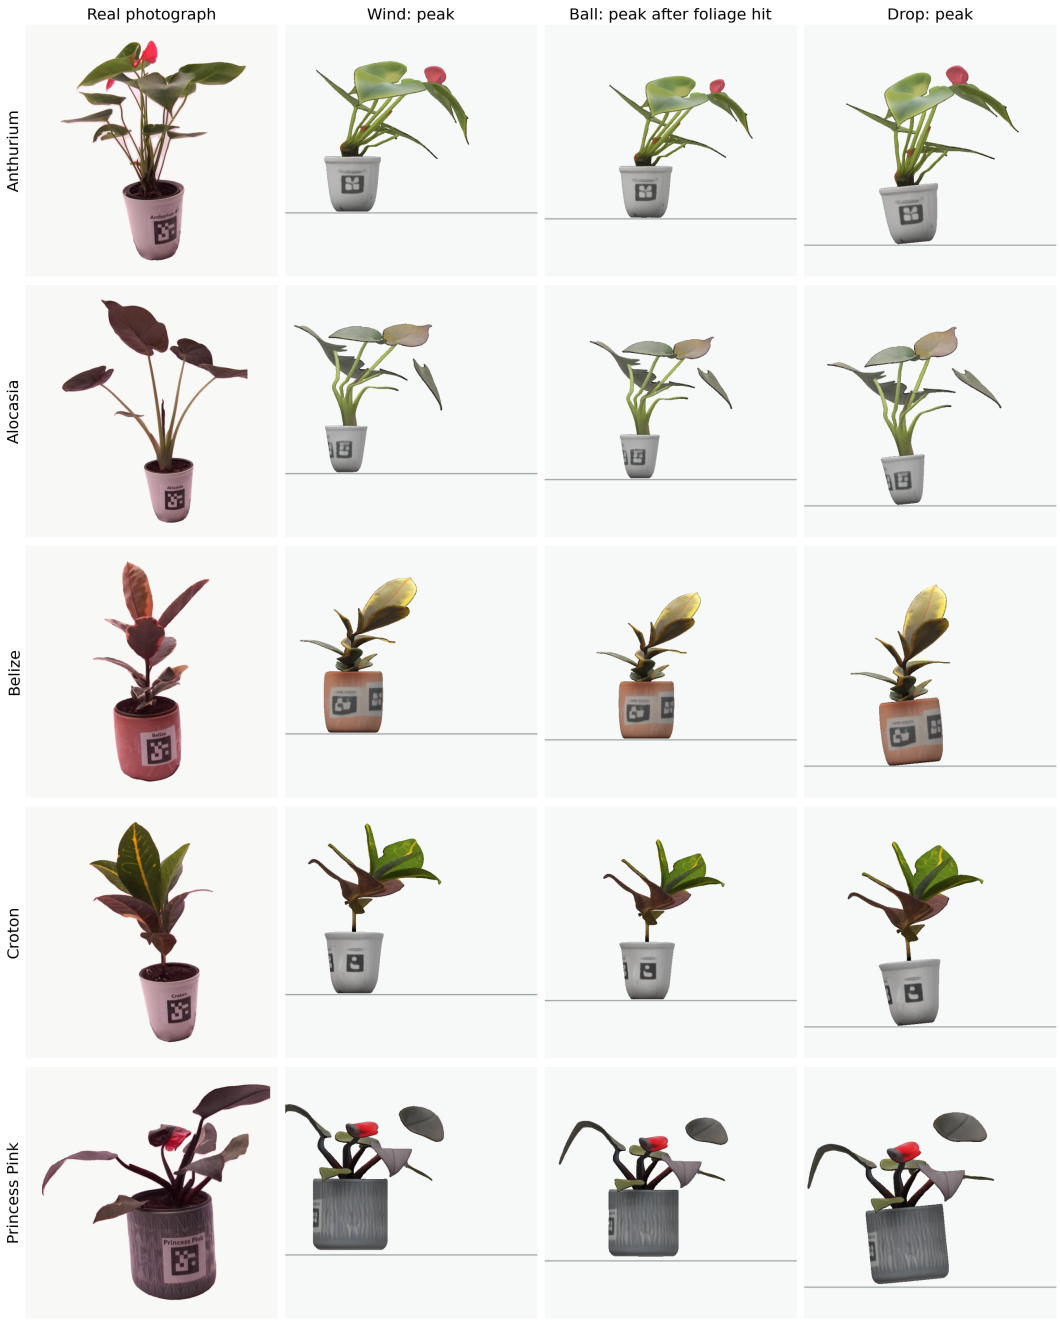

Figure 9: Peak states from the strong setting across five independently reconstructed real-photo assets. Each row contains a source photograph, peak wind bend, peak bend after the sphere contacts an upper-foliage proxy, and peak bend after the falling pot contacts the ground. Figure 10 shows one detailed contact sequence and property-status key; the full videos, per-asset keys, and saved states are in the offline supplement. These strong loads were chosen for visibility, not fitted to measured plant mechanics.

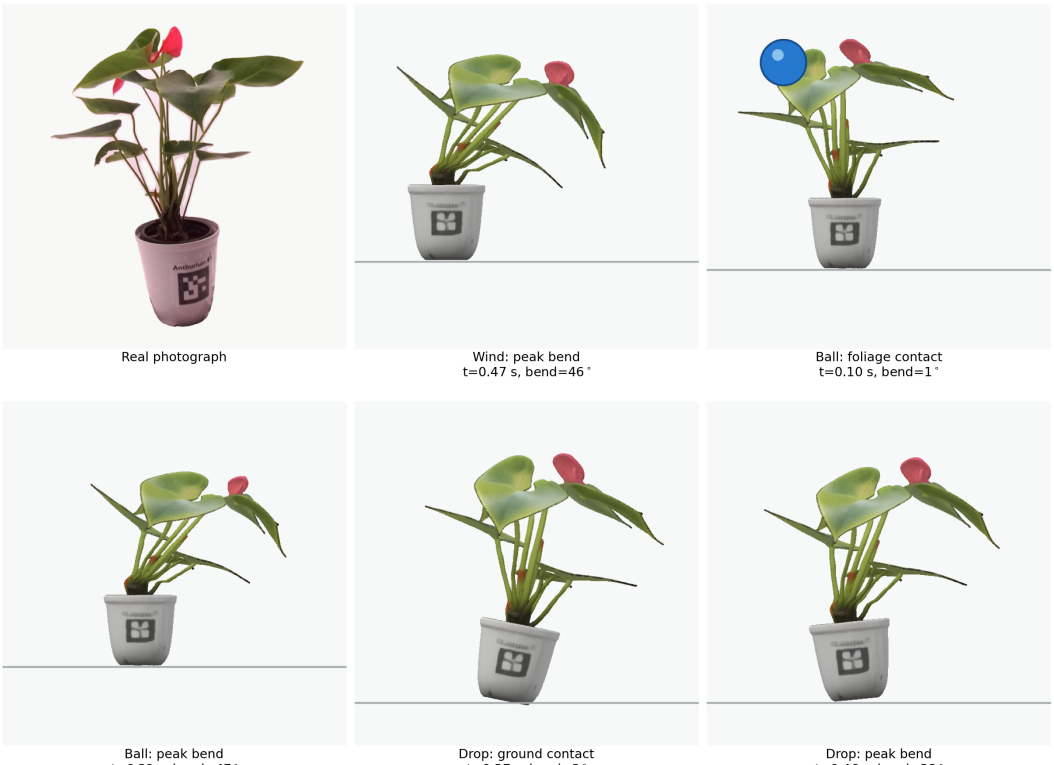

Part	ρ (kg/m^3)	E (GPa)	ν	Property status
Pot	1200	2.25	0.364	manual grade template
Soil	–	–	–	outside released property taxonomy
Stems	–	–	–	outside released property taxonomy
Leaves	–	–	–	outside released property taxonomy
Flower	–	–	–	outside released property taxonomy

Figure 10: Strong-perturbation sequence for a real-photo-derived Anthurium asset. The panels show source photograph, peak wind bend, first saved sphere–foliage contact, peak post-impact bend, first saved pot–ground contact, and peak post-contact bend. Maximum stem-proxy angles are 46°, 47°, and 28° respectively. The table enumerates property status for every named simulation part; the pot triple is a manually matched example grade, and plant/soil values are unavailable in the release. None of these numbers is a trained-model prediction for the photographed asset. Upper-foliage ellipsoids and stem capsules are MuJoCo collision proxies; the reconstructed mesh is skinned visual geometry, not a calibrated FEM solve.

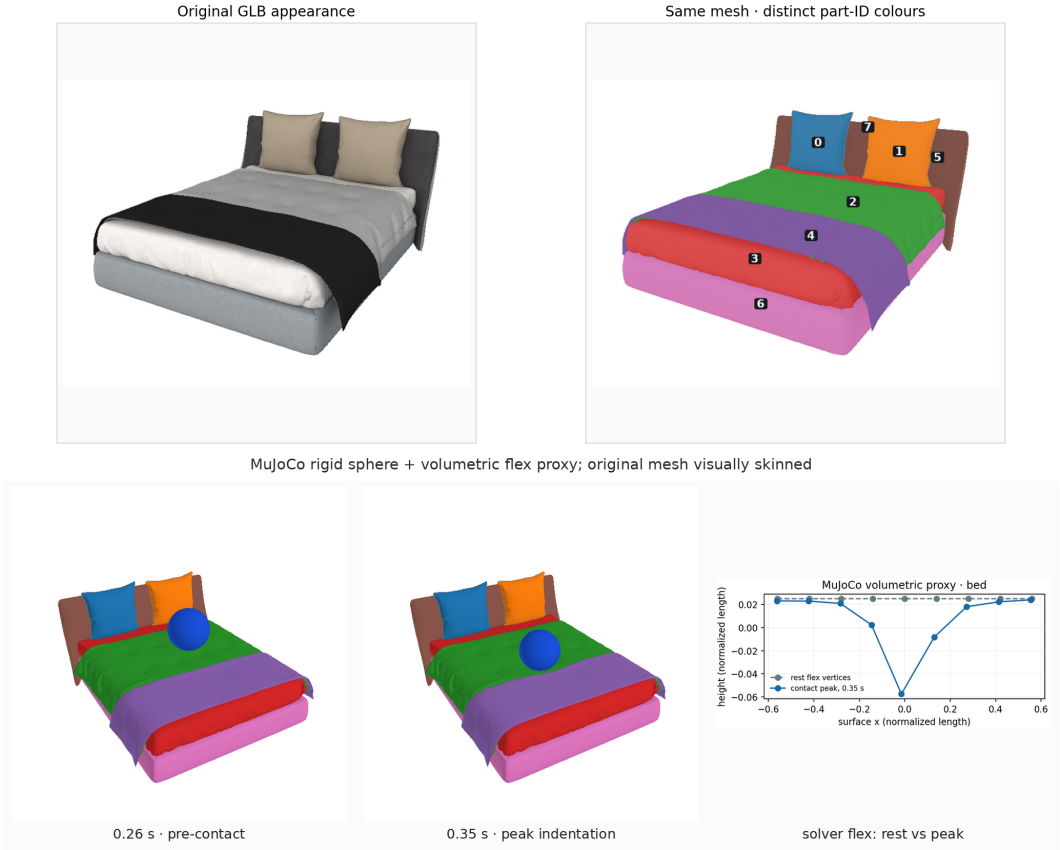

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Pillow	fabric	307.3†	0.00012†	0.3612
1	Pillow	fabric	307.3†	0.00012†	0.3612
2	Blanket	fabric	307.3†	0.00012†	0.3612
3	Mattress	fabric	307.3†	0.00012†	0.3612
4	Bed Scarf	fabric	307.3†	0.00012†	0.3612
5	Soft Surface	fabric	307.3†	0.00012†	0.3612
6	Bottom Panel	wood	560	9.215	0.3665
7	Other	fabric	307.3†	0.00012†	0.3612

† Fabric entries display the released packing-density/membrane-modulus template numbers; they are not volumetric ρ, E targets and do not set the flex proxy.

Figure 11: Held-out Bed model 43781675a0 from the PartNeXt mesh release (19,536 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables. Bottom: pre-contact and peak MuJoCo states plus the rest/peak solver-vertex cross-section for a rigid sphere contacting a pinned volumetric flex proxy, shown from a higher camera. The original GLB soft part is visually skinned to the proxy displacement. The proxy parameters are authored, not derived from the predicted properties; this is a qualitative simulator example, not measured mechanics. Contact was logged, and the nearest surface vertex moved 0.083 normalized length units at peak.

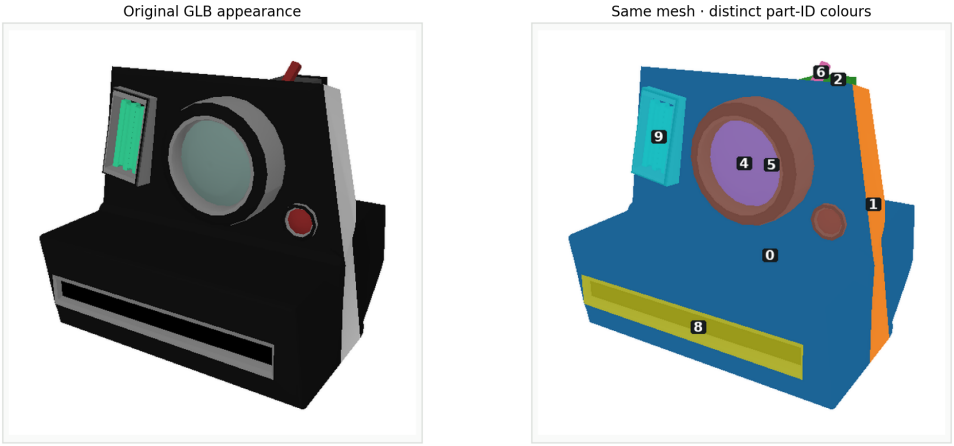

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Base	metal	2710	71	0.33
1	Casing	metal	2710	71	0.33
2	Casing	metal	2710	71	0.33
3	Lens Element	glass	2250	65.7	0.228
4	Lens Element	glass	2250	65.7	0.228
5	Casing	plastic	1200	2.25	0.3643
6	Button	plastic	1200	2.25	0.3643
7	Button	plastic	1200	2.25	0.3643
8	Other	metal	2710	71	0.33
9	Screen Body	plastic	1200	2.25	0.3643

Figure 12: Held-out Camera model c4b84f291a from the PartNeXt mesh release (1,526 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables.

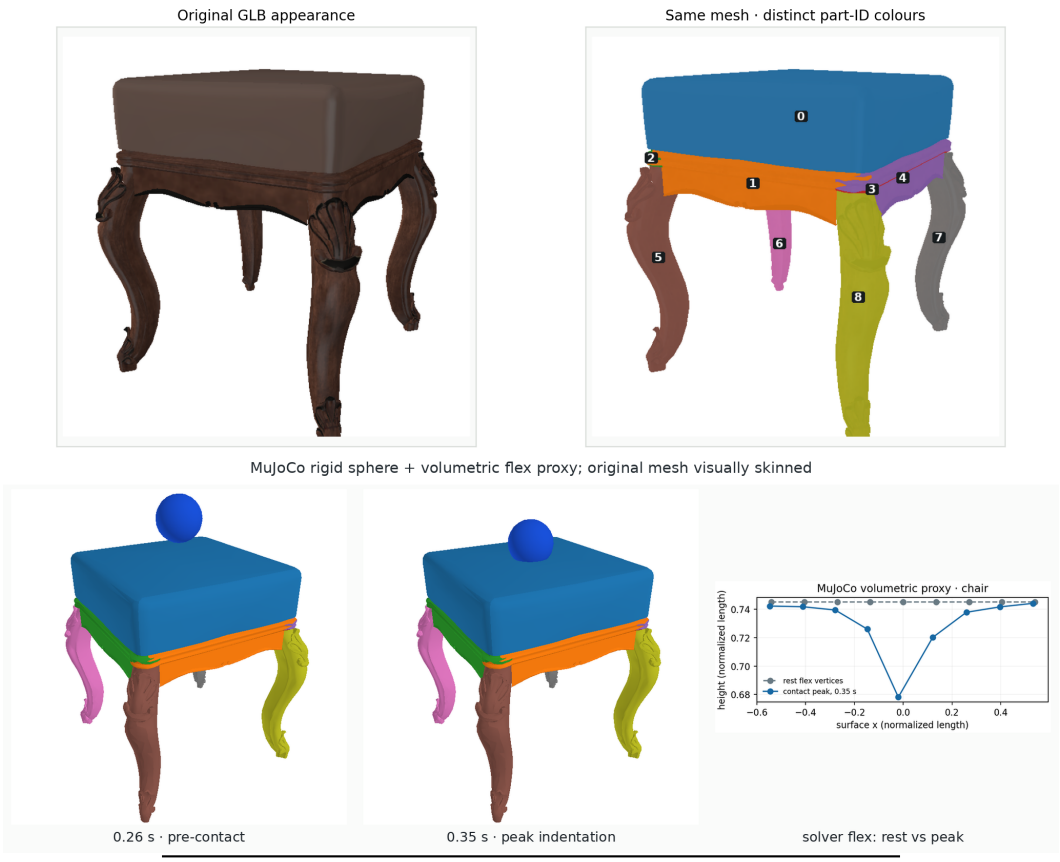

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Cushion	fabric	307.3†	0.00012†	0.3612
1	Frame	wood	560	9.215	0.3665
2	Frame	metal	8600	127.6	0.27
3	Frame	wood	560	9.215	0.3665
4	Frame	metal	8600	127.6	0.27
5	Leg	metal	8600	127.6	0.27
6	Leg	wood	560	9.215	0.3665
7	Leg	metal	8600	127.6	0.27
8	Leg	metal	8600	127.6	0.27

† Fabric entries display the released packing-density/membrane-modulus template numbers; they are not volumetric ρ, E targets and do not set the flex proxy.

Figure 13: Held-out Chair model b74eca44db from the PartNeXt mesh release (9,356 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables. Bottom: pre-contact and peak MuJoCo states plus the rest/peak solver-vertex cross-section for a rigid sphere contacting a pinned volumetric flex proxy, shown from a higher camera. The original GLB soft part is visually skinned to the proxy displacement. The proxy parameters are authored, not derived from the predicted properties; this is a qualitative simulator example, not measured mechanics. Contact was logged, and the nearest surface vertex moved 0.067 normalized length units at peak.

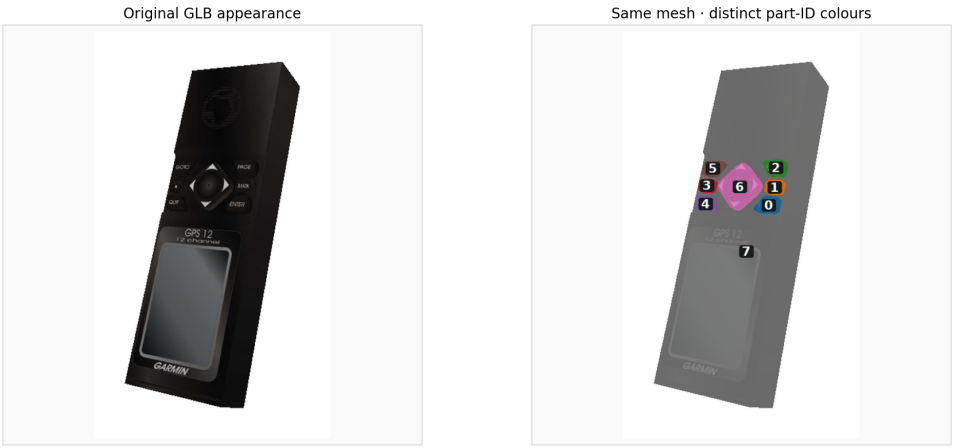

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Buttons	plastic	1200	2.25	0.3643
1	Buttons	plastic	1200	2.25	0.3643
2	Buttons	plastic	1200	2.25	0.3643
3	Buttons	plastic	1200	2.25	0.3643
4	Buttons	plastic	1200	2.25	0.3643
5	Buttons	plastic	1200	2.25	0.3643
6	Directional Pad	plastic	1200	2.25	0.3643
7	Casing	plastic	1200	2.25	0.3643

Figure 14: Held-out Controller model ae34958d91 from the PartNeXt mesh release (1,672 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables.

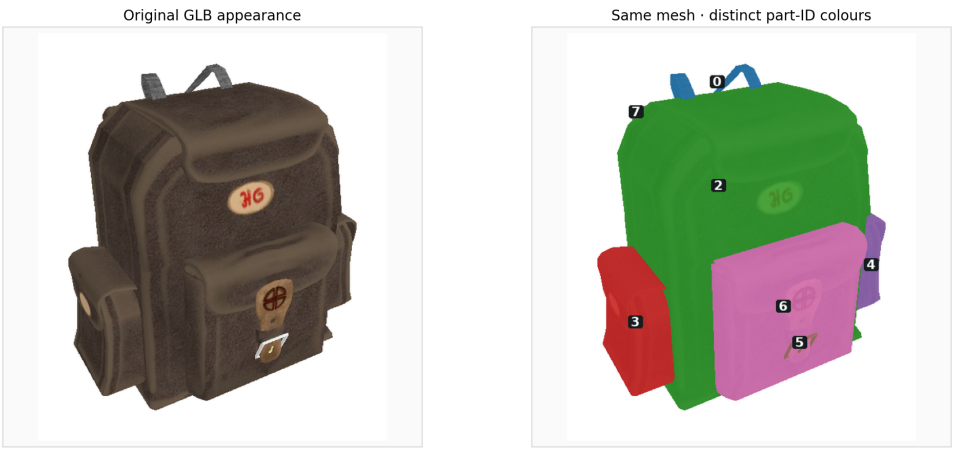

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Main Handle	rubber	1000	0.0088	0.49
1	Main Strap	fabric	307.3†	0.00012†	0.3612
2	Main Part	metal	7850	196.5	0.29
3	Pocket	fabric	307.3†	0.00012†	0.3612
4	Pocket	fabric	307.3†	0.00012†	0.3612
5	Lock	metal	7850	196.5	0.29
6	Pocket	fabric	307.3†	0.00012†	0.3612
7	Main Strap	fabric	307.3†	0.00012†	0.3612

† Fabric entries display the released packing-density/membrane-modulus template numbers; they are not volumetric ρ, E targets and do not set the flex proxy.

Figure 15: Held-out Handbag model e8fe8902cd from the PartNeXt mesh release (6,062 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables.

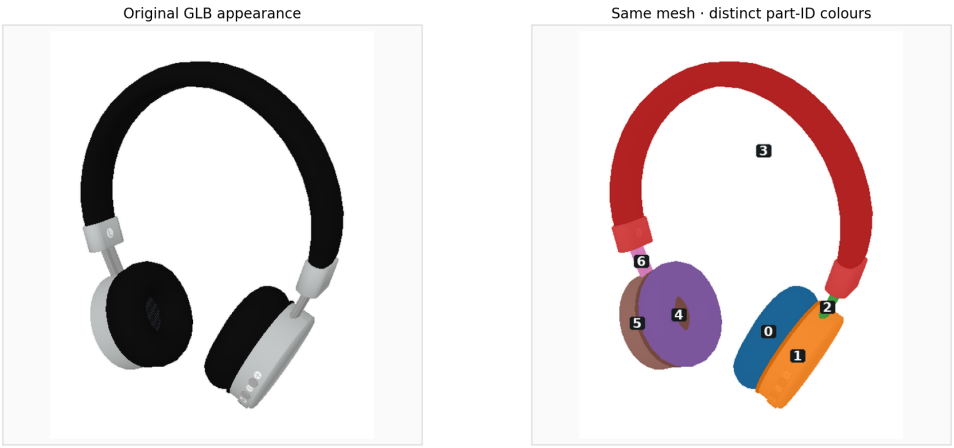

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Ear Pad	fabric	307.3†	0.00012†	0.3612
1	Ear Cup Frame	metal	7850	203	0.29
2	Ear Cup Connector	plastic	1200	2.25	0.3643
3	Headband	metal	7850	203	0.29
4	Ear Pad	fabric	307.3†	0.00012†	0.3612
5	Ear Cup Frame	metal	7850	203	0.29
6	Ear Cup Connector	plastic	1200	2.25	0.3643

† Fabric entries display the released packing-density/membrane-modulus template numbers; they are not volumetric ρ, E targets and do not set the flex proxy.

Figure 16: Held-out Headphone model 9f469c795d from the PartNeXt mesh release (4,627 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables.

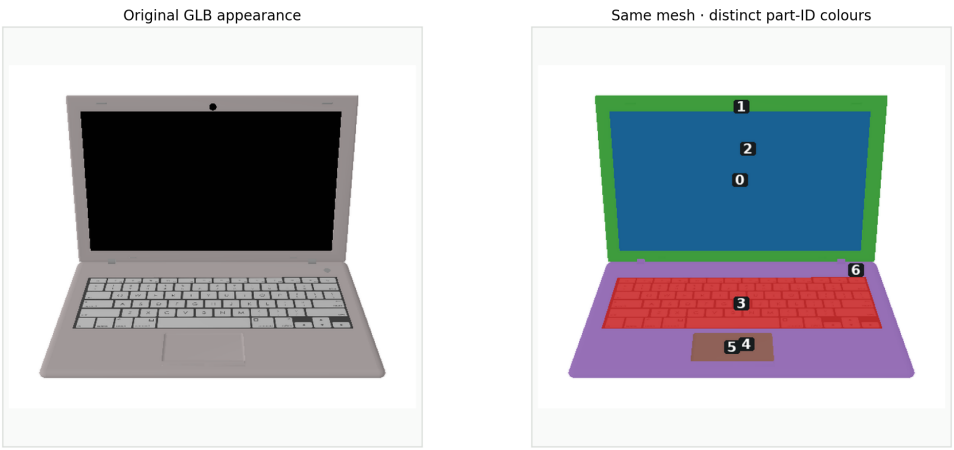

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Screen	glass	2250	65.7	0.228
1	Camera	plastic	1200	2.25	0.3643
2	Casing	metal	2710	71	0.33
3	Backplate	metal	7850	196.5	0.29
4	Casing	metal	2710	71	0.33
5	Touchpad	plastic	1200	2.25	0.3643
6	Buttons	plastic	1200	2.25	0.3643

Figure 17: Held-out Laptop Computer model 9b45e7a38d from the PartNeXt mesh release (10,014 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables.

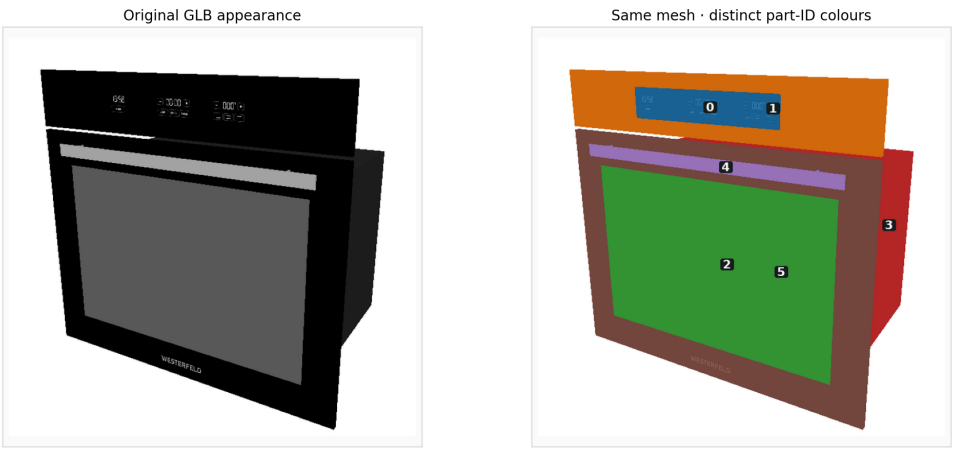

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Display Screen	plastic	1200	2.25	0.3643
1	Control Panel Backplate	metal	7800	196.5	0.29
2	Door Glass	glass	2230	62.6	0.228
3	Casing	metal	7800	196.5	0.29
4	Door Handle	metal	7800	196.5	0.29
5	Door Frame	metal	7850	203	0.29

Figure 18: Held-out Microwave Oven model be59c580bd from the PartNeXt mesh release (22,670 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables.

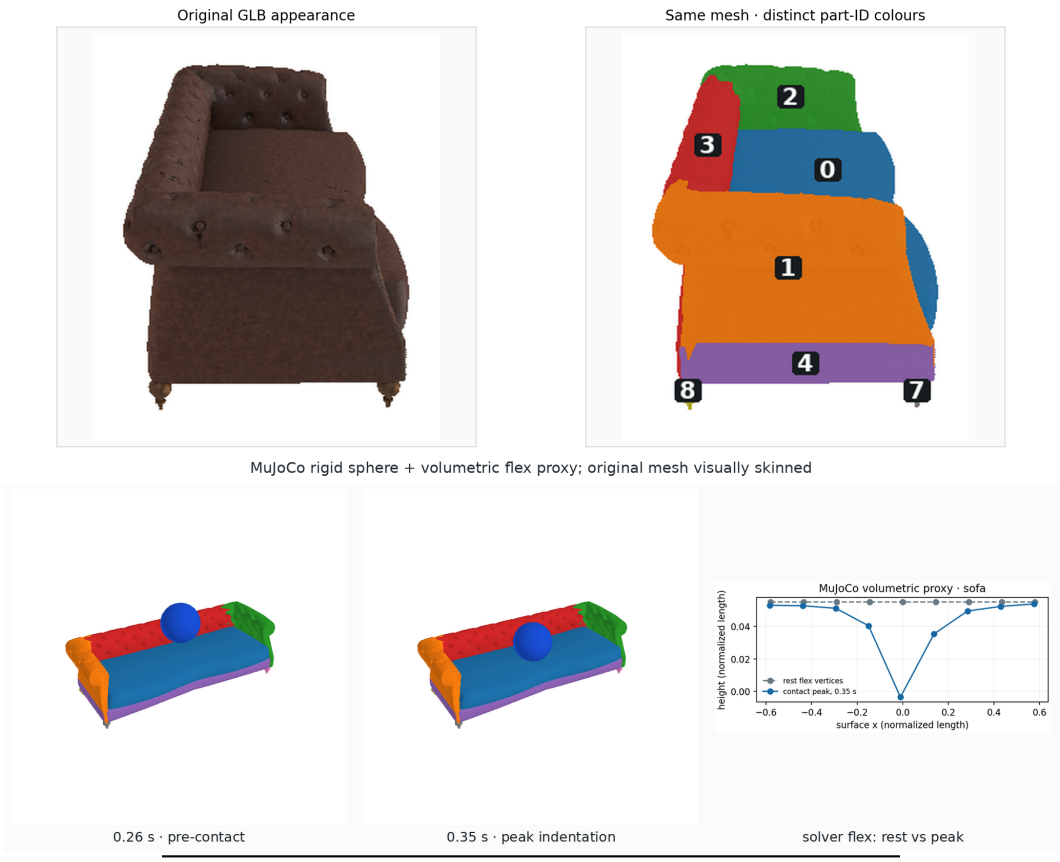

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Seat Cushion	fabric	307.3†	0.00012†	0.3612
1	Armrest Cushion	fabric	307.3†	0.00012†	0.3612
2	Armrest Cushion	fabric	307.3†	0.00012†	0.3612
3	Backrest Cushion	fabric	307.3†	0.00012†	0.3612
4	Seat Cushion Support Box	wood	560	9.215	0.3665
5	Leg	metal	8600	110	0.33
6	Leg	wood	560	9.215	0.3665
7	Leg	wood	560	9.215	0.3665
8	Leg	wood	560	9.215	0.3665

† Fabric entries display the released packing-density/membrane-modulus template numbers; they are not volumetric ρ, E targets and do not set the flex proxy.

Figure 19: Held-out Sofa model 58d54c95c3 from the PartNeXt mesh release (19,699 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables. Bottom: pre-contact and peak MuJoCo states plus the rest/peak solver-vertex cross-section for a rigid sphere contacting a pinned volumetric flex proxy, shown from a higher camera. The original GLB soft part is visually skinned to the proxy displacement. The proxy parameters are authored, not derived from the predicted properties; this is a qualitative simulator example, not measured mechanics. Contact was logged, and the nearest surface vertex moved 0.059 normalized length units at peak.

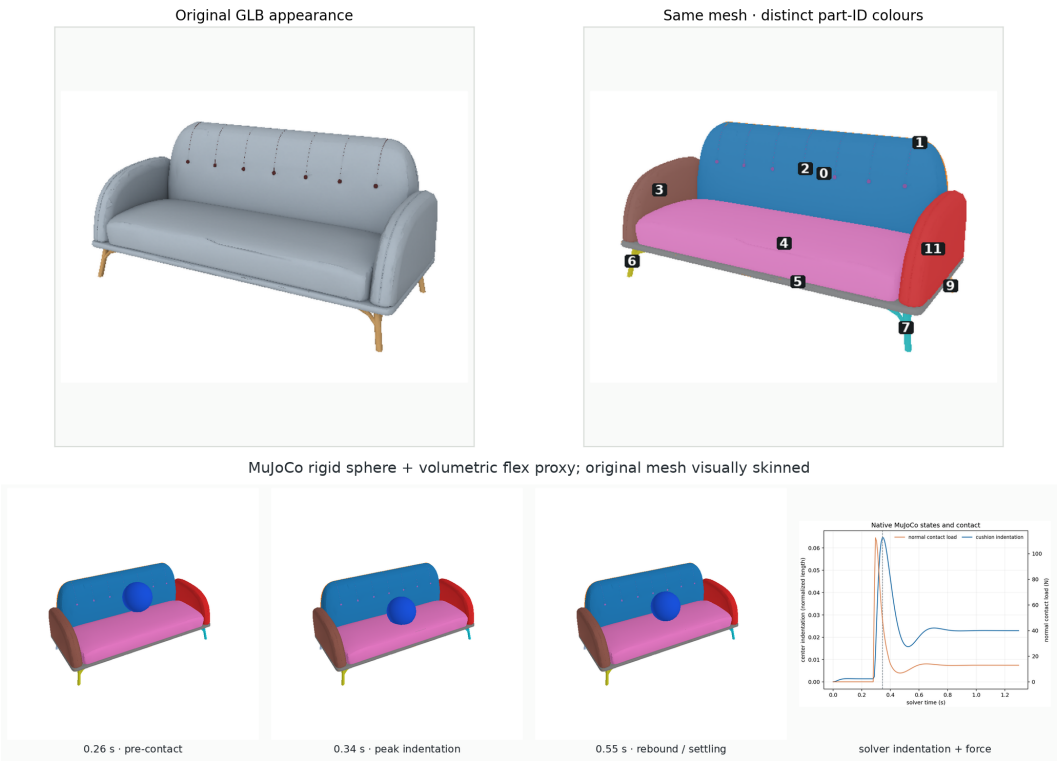

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Backrest Cushion	fabric	307.3†	0.00012†	0.3612
1	Backrest Frame	fabric	307.3†	0.00012†	0.3612
10	Other	fabric	307.3†	0.00012†	0.3612
11	Armrest Support Box	wood	560	9.215	0.3665
2	Other	fabric	307.3†	0.00012†	0.3612
3	Armrest Support Box	wood	560	9.215	0.3665
4	Seat Cushion	fabric	307.3†	0.00012†	0.3612
5	Seat Cushion Support Box	wood	560	9.215	0.3665
6	Leg	wood	560	9.215	0.3665
7	Leg	wood	560	9.215	0.3665
8	Leg	wood	560	9.215	0.3665
9	Leg	wood	560	9.215	0.3665

† Fabric entries display the released packing-density/membrane-modulus template numbers; they are not volumetric ρ, E targets and do not set the flex proxy.

Figure 20: Held-out Sofa model 64f9de49a7 from the PartNeXt mesh release (20,056 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables. Bottom: the rigid sphere strikes part ID 4 (seat cushion); the three rendered states come from saved MuJoCo pinned-volumetric-flex vertex positions, followed by the solver's center indentation and normal contact-load curves. Original GLB triangles are skinned for display; the grid flex is the actual collision body. Parameters are authored rather than derived from the material predictor; the episode is not calibrated real-world mechanics. Peak center displacement is 0.065 normalized length units.



ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Surface Panel	wood	560	9.215	0.3665
1	Table Frame	wood	560	9.215	0.3665
2	Leg	wood	560	9.215	0.3665
3	Leg	wood	560	9.215	0.3665
4	Leg	wood	560	9.215	0.3665
5	Leg	wood	560	9.215	0.3665
6	Other	wood	560	9.215	0.3665

Figure 21: Held-out Table model 1338705a26 from the PartNeXt mesh release (32,619 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables.

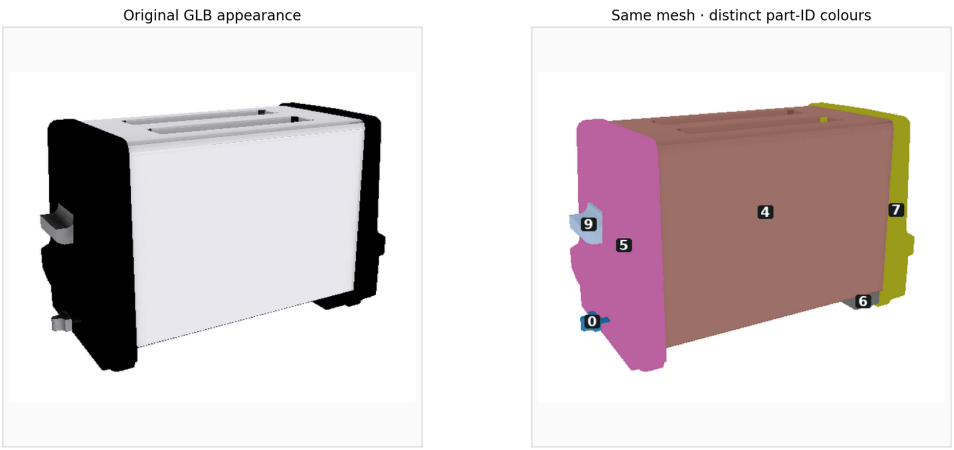

ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Lever Body	metal	7850	196.5	0.29
1	Slide Rail	plastic	1200	2.25	0.3643
10	Inner Shell	metal	7850	196.5	0.29
2	Horizontal Bar	metal	7850	203	0.29
3	Horizontal Bar	metal	7850	203	0.29
4	Casing	metal	7850	196.5	0.29
5	Bread Slot Side Wall	metal	8600	127.6	0.27
6	Support Foot	rubber	1015	0.000646	0.49
7	Bread Slot Side Wall	metal	8600	127.6	0.27
8	Slide Rail	plastic	1200	2.25	0.3643
9	Lever Body	plastic	1200	2.25	0.3643

Figure 22: Held-out Toaster model 68a31a8eb0 from the PartNeXt mesh release (22,628 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables.



ID	Part	Predicted grade	ρ (kg/m^3)	E (GPa)	ν
0	Frame	metal	7850	203	0.29
1	Door Frame	metal	7850	203	0.29
10	foot	metal	7850	203	0.29
2	Inner Drum	metal	7850	196.5	0.29
3	Drum Shaft	metal	7850	196.5	0.29
5	Frame	metal	7850	203	0.29
6	Knob	metal	7850	196.5	0.29
7	foot	metal	7850	203	0.29
8	foot	metal	7850	203	0.29
9	foot	metal	7850	203	0.29

Figure 23: Held-out Washing Machine model 7e088f1294 from the PartNeXt mesh release (50,308 original GLB triangles). Top left: original mesh appearance. Top right: each part has a distinct ID colour keyed to the table swatch; material grades and properties are frozen-model predictions. Values are predictions of released material priors, not measured mechanics. All parts, including invisible ones, are listed; released labels and errors appear in the complete 88-asset tables. This source GLB contains grayscale vertex colours but no embedded texture image or UVs; the left panel renders those source vertex colours.

L ADDITIONAL HELD-OUT COLLISION TARGETS

Figures 24–26 apply the same saved-state renderer to three further sphere–soft-proxy targets: both pillows of the bed in Figure 11, and the seat cushion of the second held-out sofa in Figure 19. The source meshes are unchanged and their complete part-property tables appear with those case studies. The two pillow episodes share the bed GLB but use different original part IDs and separately placed collision proxies. Each target also has light and heavy load variants in the offline video supplement. These examples increase visual and contact diversity; they do not provide a measured-mechanics benchmark or demonstrate a gain from the material predictor.



Figure 24: Additional contact target on held-out Bed model 43781675a0: left pillow (original part ID 0). The first two panels render the original PartNeXt GLB triangles in distinct part-ID colours before contact and at peak indentation. The graph shows the native MuJoCo pinned-flex surface vertices at rest and peak; the center vertex moves 0.057 normalized length units, with 484 solver steps registering sphere–flex contact. The original GLB is visually skinned, while the grid flex is the collision body. All proxy parameters and loads are authored, not predicted physical properties. The offline supplement supplies original-appearance and part-coloured videos, scene XML, states, and the native solver render.



Figure 25: Additional contact target on held-out Bed model 43781675a0: right pillow (original part ID 1). The first two panels render the original PartNeXt GLB triangles in distinct part-ID colours before contact and at peak indentation. The graph shows the native MuJoCo pinned-flex surface vertices at rest and peak; the center vertex moves 0.057 normalized length units, with 484 solver steps registering sphere–flex contact. The original GLB is visually skinned, while the grid flex is the collision body. All proxy parameters and loads are authored, not predicted physical properties. The offline supplement supplies original-appearance and part-coloured videos, scene XML, states, and the native solver render.

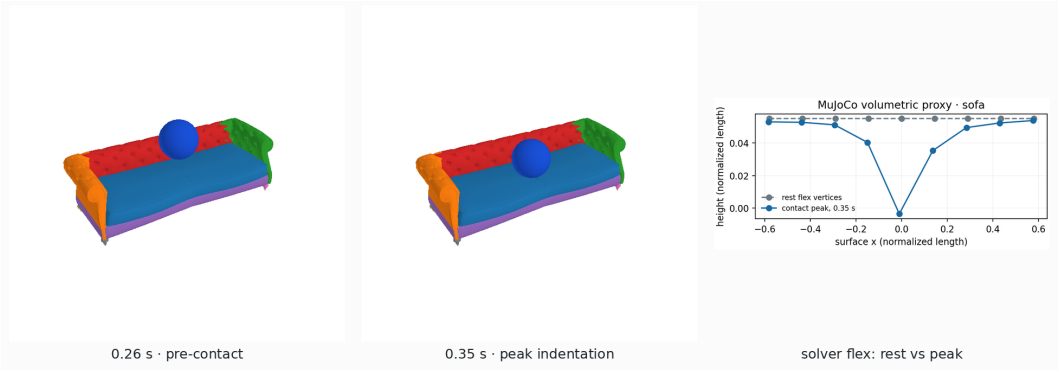

Figure 26: Additional contact target on held-out Sofa model 58d54c95c3: seat cushion (original part ID 0). The first two panels render the original PartNeXt GLB triangles in distinct part-ID colours before contact and at peak indentation. The graph shows the native MuJoCo pinned-flex surface vertices at rest and peak; the center vertex moves 0.059 normalized length units, with 508 solver steps registering sphere–flex contact. The original GLB is visually skinned, while the grid flex is the collision body. All proxy parameters and loads are authored, not predicted physical properties. The offline supplement supplies original-appearance and part-coloured videos, scene XML, states, and the native solver render.

M RIGID-BLOCK CONTACT ON HELD-OUT SOFT FURNITURE

To expose a broader contact patch than the sphere examples, we also release eight rigid-block/flex episodes on four of the same held-out PartNeXt GLBs: a bed mattress/blanket, a chair cushion, and two sofa seats. Medium and strong settings change the authored block mass from 3.0 to 4.5 solver mass units and its normalized drop height from 0.68 to 0.78. In every case, a free rigid block contacts a pinned tetrahedral flex in MuJoCo; the saved vertex states are replayed for a native solver render. The original GLB retains its stored appearance or part-ID colours and is visually skinned to the flex displacement. It does not supply the collision elements. Table 23 reports the nearest center flex-surface vertex, and Figure 31 shows the solver traces. All four strong settings produce larger center displacement than their medium counterparts. These results demonstrate functioning collision-state generation and visual transfer to original meshes, but do not test whether the predicted (ρ, E, ν) improves mechanics: the flex parameters, block masses, and boundary conditions were authored. The four source figures cited in the captions contain all part-level predictions, including fabric applicability masks and template conventions. The accompanying web cards provide synchronized part-coloured, original-appearance, and native solver videos together with XML and saved states.

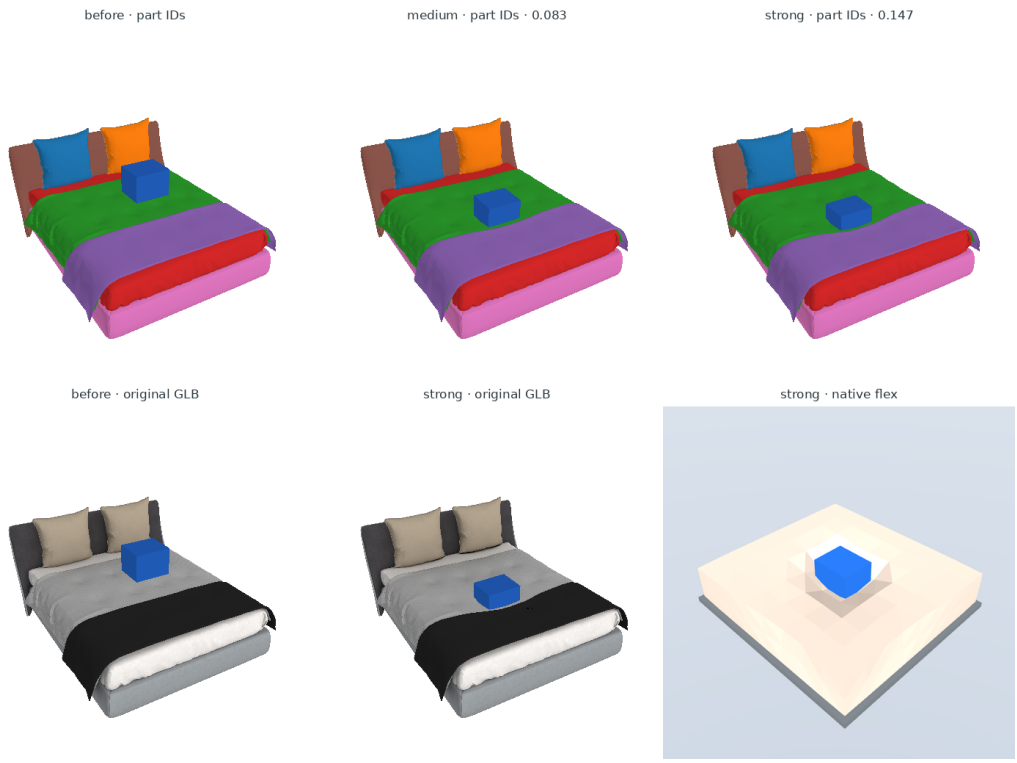

Figure 27: Rigid-block contact on the held-out Bed mattress/blanket (model 43781675a0; visual soft part IDs 2, 3, 4). The upper row renders the original GLB triangles with distinct part-ID colours before impact and at the two saved peak-contact states. The lower row retains the source appearance before and at strong contact, then shows the native MuJoCo rigid block and tetrahedral flex proxy at the strong peak. Medium and strong center-vertex displacements are 0.083 and 0.147 normalized length units, with 477 and 463 solver contact steps. The original GLB is visually skinned from the proxy motion; its triangles are not the collision mesh. Flex and load parameters are authored, not set from predicted material properties. Every part's predicted property appears in Figure 11 and the offline supplement.

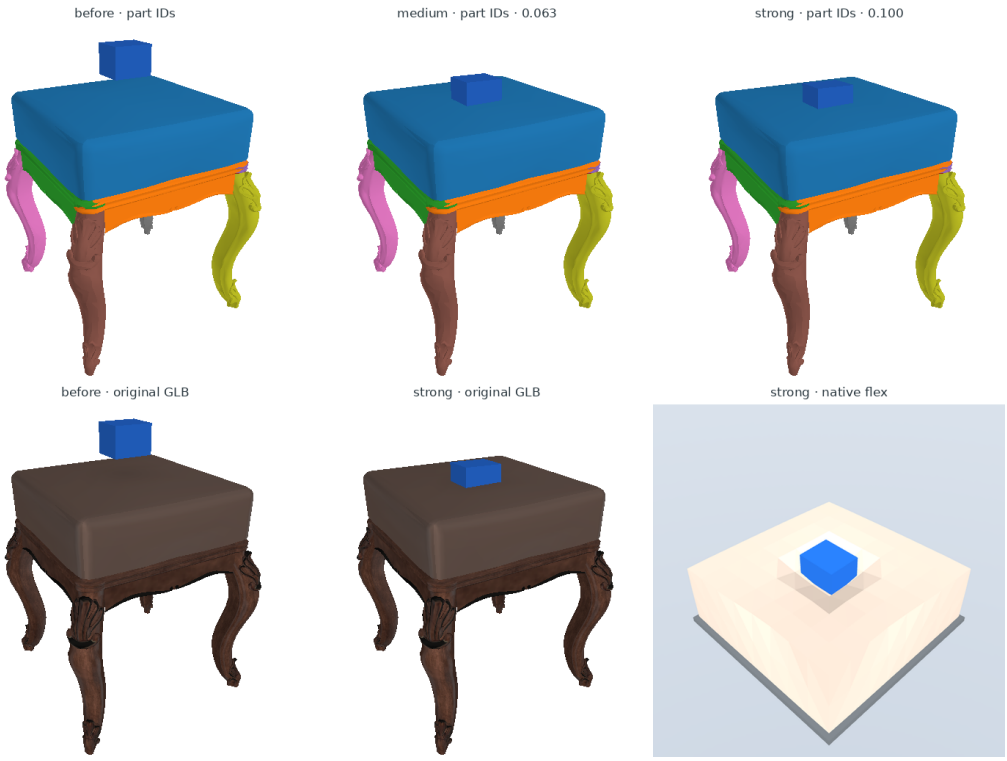

Figure 28: Rigid-block contact on the held-out Chair cushion (model b74eca44db; visual soft part IDs 0). The upper row renders the original GLB triangles with distinct part-ID colours before impact and at the two saved peak-contact states. The lower row retains the source appearance before and at strong contact, then shows the native MuJoCo rigid block and tetrahedral flex proxy at the strong peak. Medium and strong center-vertex displacements are 0.063 and 0.100 normalized length units, with 477 and 463 solver contact steps. The original GLB is visually skinned from the proxy motion; its triangles are not the collision mesh. Flex and load parameters are authored, not set from predicted material properties. Every part's predicted property appears in Figure 13 and the offline supplement.

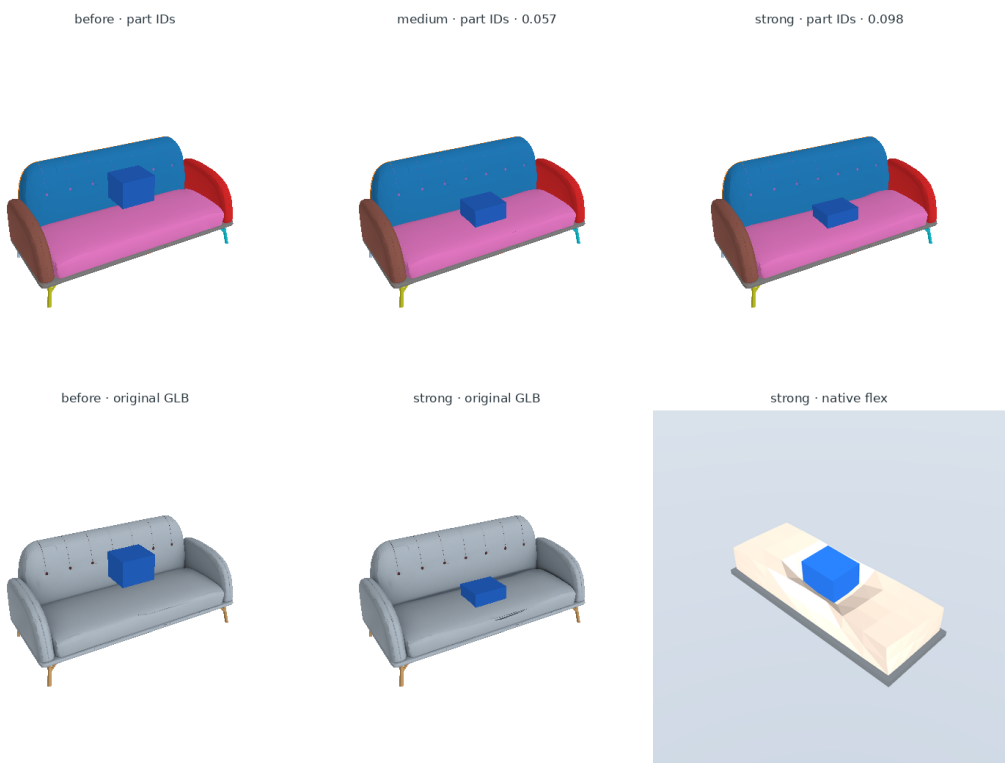

Figure 29: Rigid-block contact on the held-out First sofa seat (model 64f9de49a7; visual soft part IDs 4). The upper row renders the original GLB triangles with distinct part-ID colours before impact and at the two saved peak-contact states. The lower row retains the source appearance before and at strong contact, then shows the native MuJoCo rigid block and tetrahedral flex proxy at the strong peak. Medium and strong center-vertex displacements are 0.057 and 0.098 normalized length units, with 477 and 463 solver contact steps. The original GLB is visually skinned from the proxy motion; its triangles are not the collision mesh. Flex and load parameters are authored, not set from predicted material properties. Every part's predicted property appears in Figure 20 and the offline supplement.

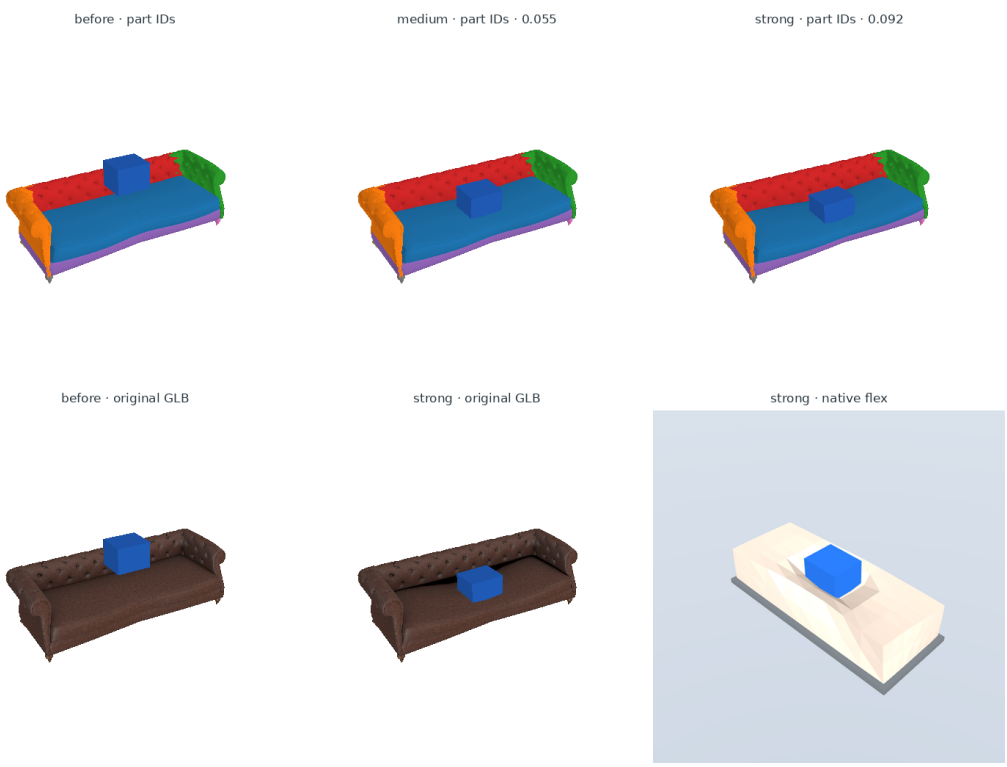

Figure 30: Rigid-block contact on the held-out Second sofa seat (model 58d54c95c3; visual soft part IDs 0). The upper row renders the original GLB triangles with distinct part-ID colours before impact and at the two saved peak-contact states. The lower row retains the source appearance before and at strong contact, then shows the native MuJoCo rigid block and tetrahedral flex proxy at the strong peak. Medium and strong center-vertex displacements are 0.055 and 0.092 normalized length units, with 477 and 463 solver contact steps. The original GLB is visually skinned from the proxy motion; its triangles are not the collision mesh. Flex and load parameters are authored, not set from predicted material properties. Every part's predicted property appears in Figure 19 and the offline supplement.

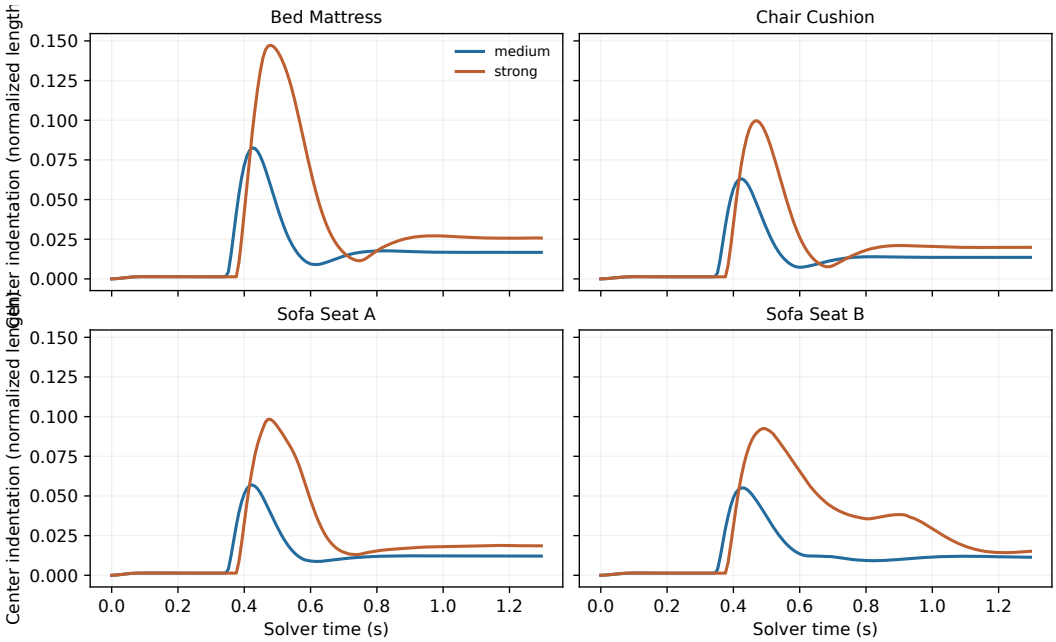

Figure 31: Saved MuJoCo center-vertex displacement traces for the four original-GLB targets under two authored rigid-block loads. The plotted values come from the pinned volumetric-flex solver, not from image-based deformation estimates or measured physical tests.

Target	Visual part ID(s)	Medium	Strong	Contact steps
Bed mattress/blanket	2, 3, 4	0.083	0.147	477 / 463
Chair cushion	0	0.063	0.100	477 / 463
First sofa seat	4	0.057	0.098	477 / 463
Second sofa seat	0	0.055	0.092	477 / 463

Table 23: Peak downward displacement of the nearest center flex-surface vertex, in normalized model-length units. The final column reports medium/strong MuJoCo solver steps with contact. These values describe authored proxy scenes, not measured material-prediction accuracy.

M.1 Saved contact-to-later-state sequences

Figures 32–35 expand four strong-setting rigid-block episodes. Each column combines the same saved time point in the original GLB appearance, per-part colour rendering, and native MuJoCo geometry; lower graphs expose the center flex-vertex trajectory and contact force. Video frame numbers, exact saved-state indices and timestamps are retained in the accompanying process manifest. Native video replay differed from saved flex vertices by less than 3.1×10^{-8} normalized units in every case. The medium and strong runs alter both mass and drop height; only Figure 39 isolates rigid density at a single bed location.

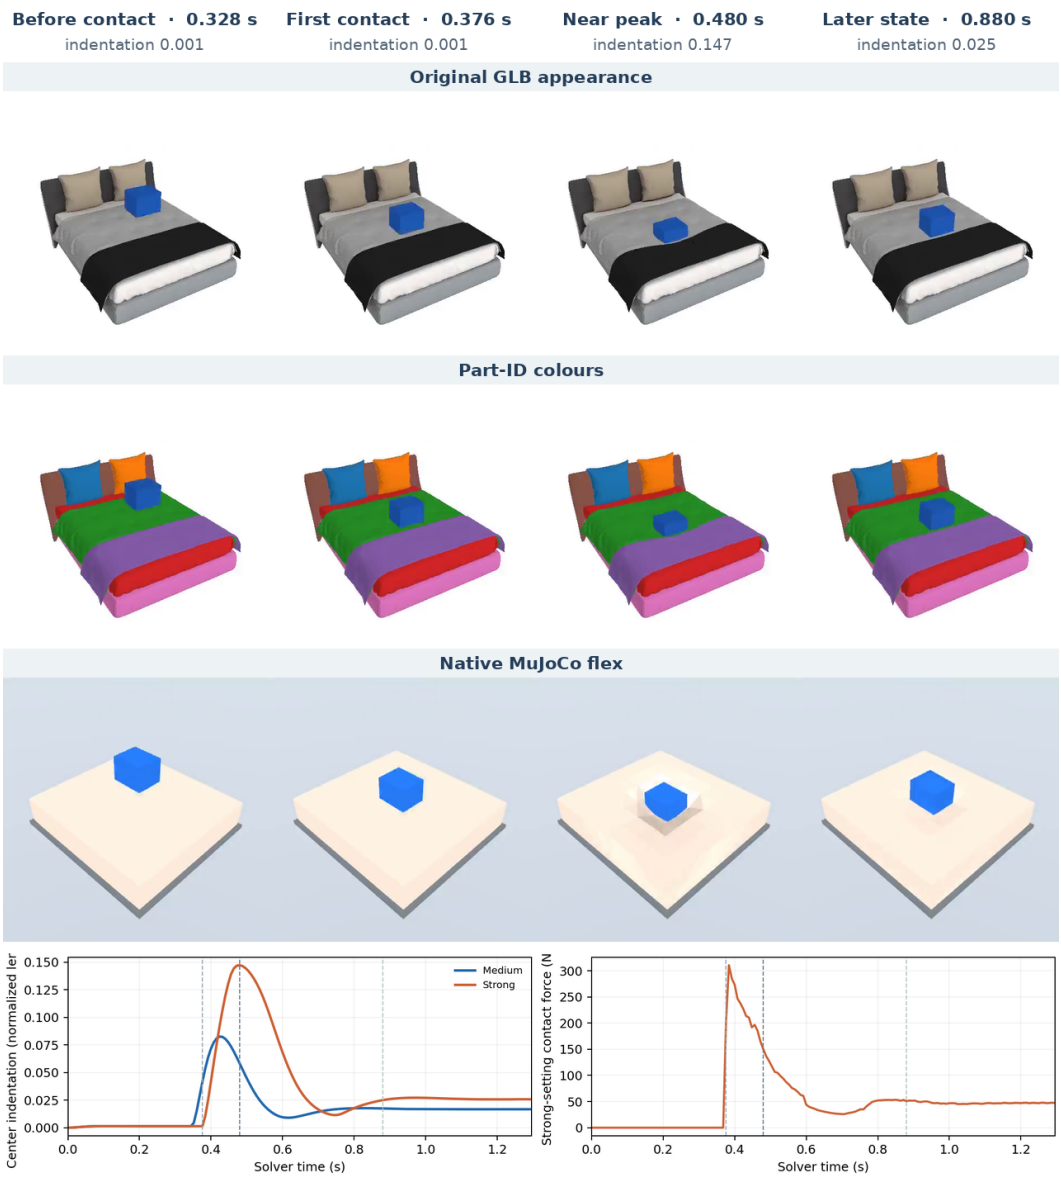

Figure 32: Saved rigid-block contact process for the held-out bed mattress/blanket (model 43781675a0). Columns show the same saved-state indices before contact, at first recorded contact, nearest the solver's peak center-vertex indentation, and later in the response. Rows show original GLB appearance, distinct part-ID colours, and native MuJoCo block/flex geometry. Times are 0.328, 0.376, 0.480, 0.880 s, with center indentations 0.001, 0.001, 0.147, 0.025 normalized length units. Lower plots show center indentation under medium and strong authored loads and the strong-run normal contact force (peak 310.7 N). Medium and strong settings change both rigid mass and drop height, so this is a process comparison rather than an isolated stiffness test. The original GLB triangles are visually skinned from saved proxy motion; the native solver uses a coarse pinned volumetric flex, and predicted material priors do not set its parameters. All part-property predictions appear in Figure 11.

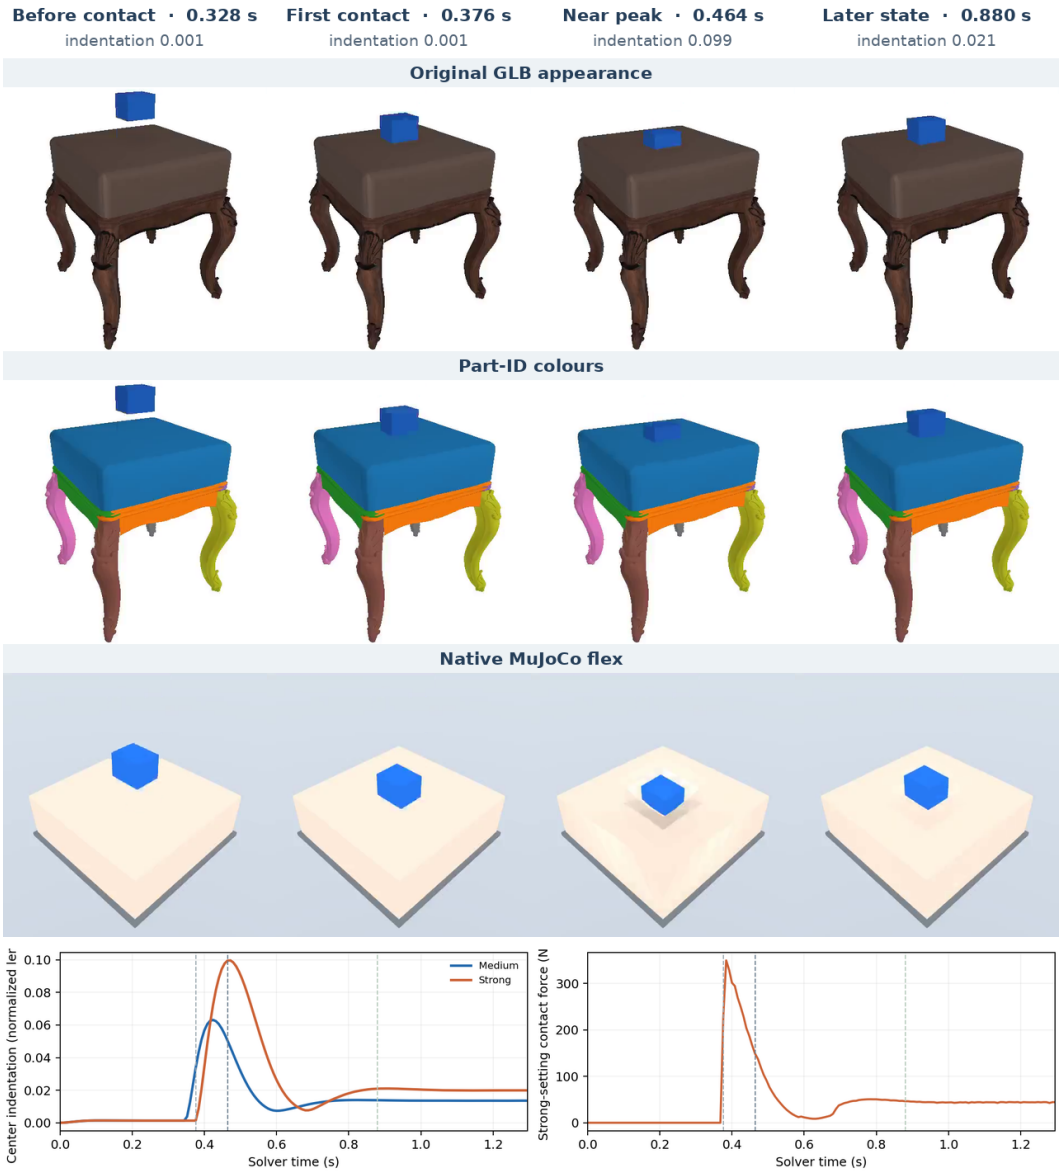

Figure 33: Saved rigid-block contact process for the held-out chair cushion (model b74eca44db). Columns show the same saved-state indices before contact, at first recorded contact, nearest the solver's peak center-vertex indentation, and later in the response. Rows show original GLB appearance, distinct part-ID colours, and native MuJoCo block/flex geometry. Times are 0.328, 0.376, 0.464, 0.880 s, with center indentations 0.001, 0.001, 0.099, 0.021 normalized length units. Lower plots show center indentation under medium and strong authored loads and the strong-run normal contact force (peak 349.4 N). Medium and strong settings change both rigid mass and drop height, so this is a process comparison rather than an isolated stiffness test. The original GLB triangles are visually skinned from saved proxy motion; the native solver uses a coarse pinned volumetric flex, and predicted material priors do not set its parameters. All part-property predictions appear in Figure 13.

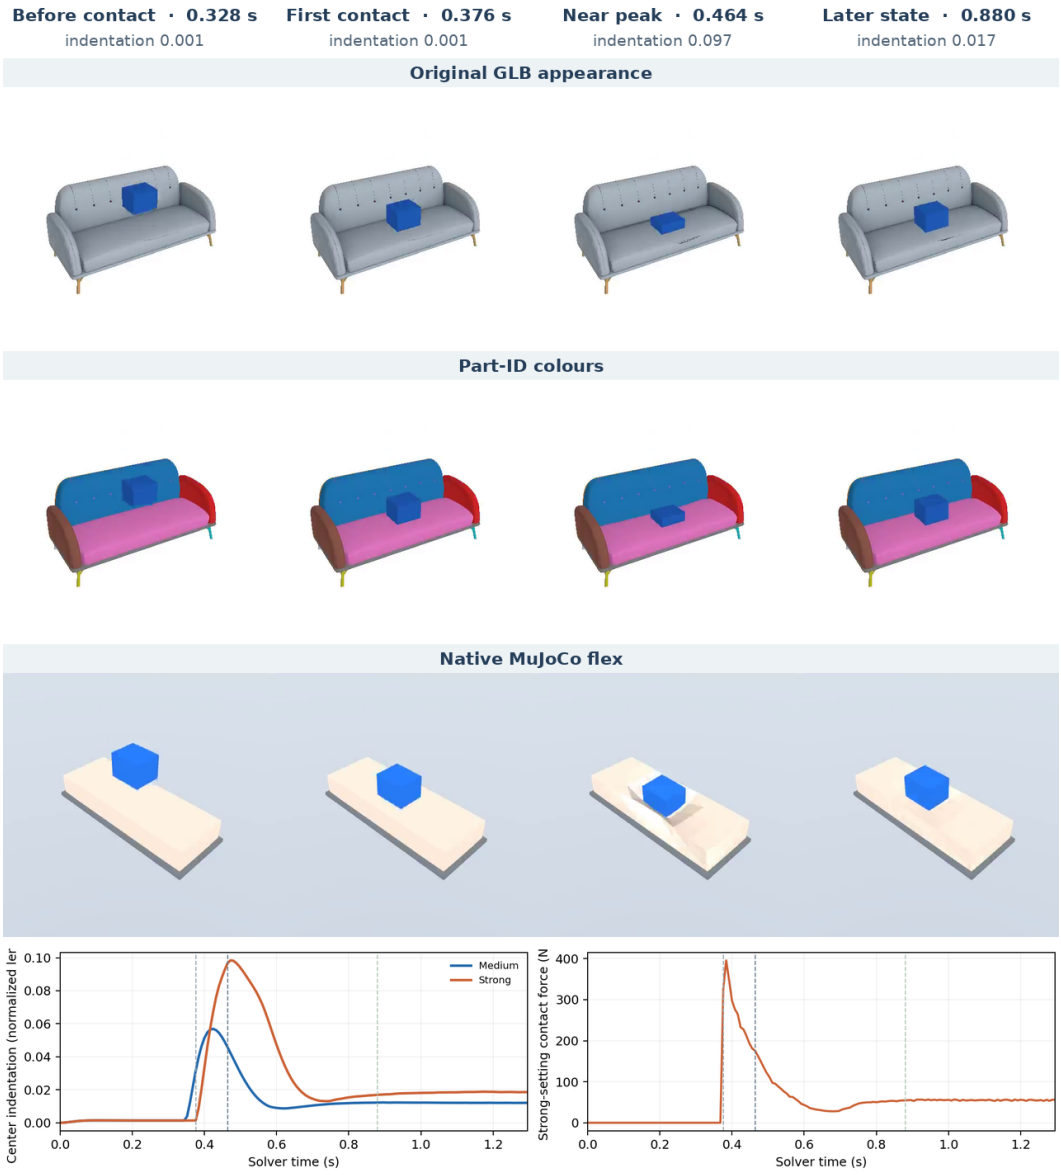

Figure 34: Saved rigid-block contact process for the held-out first sofa seat (model 64f9de49a7). Columns show the same saved-state indices before contact, at first recorded contact, nearest the solver's peak center-vertex indentation, and later in the response. Rows show original GLB appearance, distinct part-ID colours, and native MuJoCo block/flex geometry. Times are 0.328, 0.376, 0.464, 0.880 s, with center indentations 0.001, 0.001, 0.097, 0.017 normalized length units. Lower plots show center indentation under medium and strong authored loads and the strong-run normal contact force (peak 395.8 N). Medium and strong settings change both rigid mass and drop height, so this is a process comparison rather than an isolated stiffness test. The original GLB triangles are visually skinned from saved proxy motion; the native solver uses a coarse pinned volumetric flex, and predicted material priors do not set its parameters. All part-property predictions appear in Figure 20.

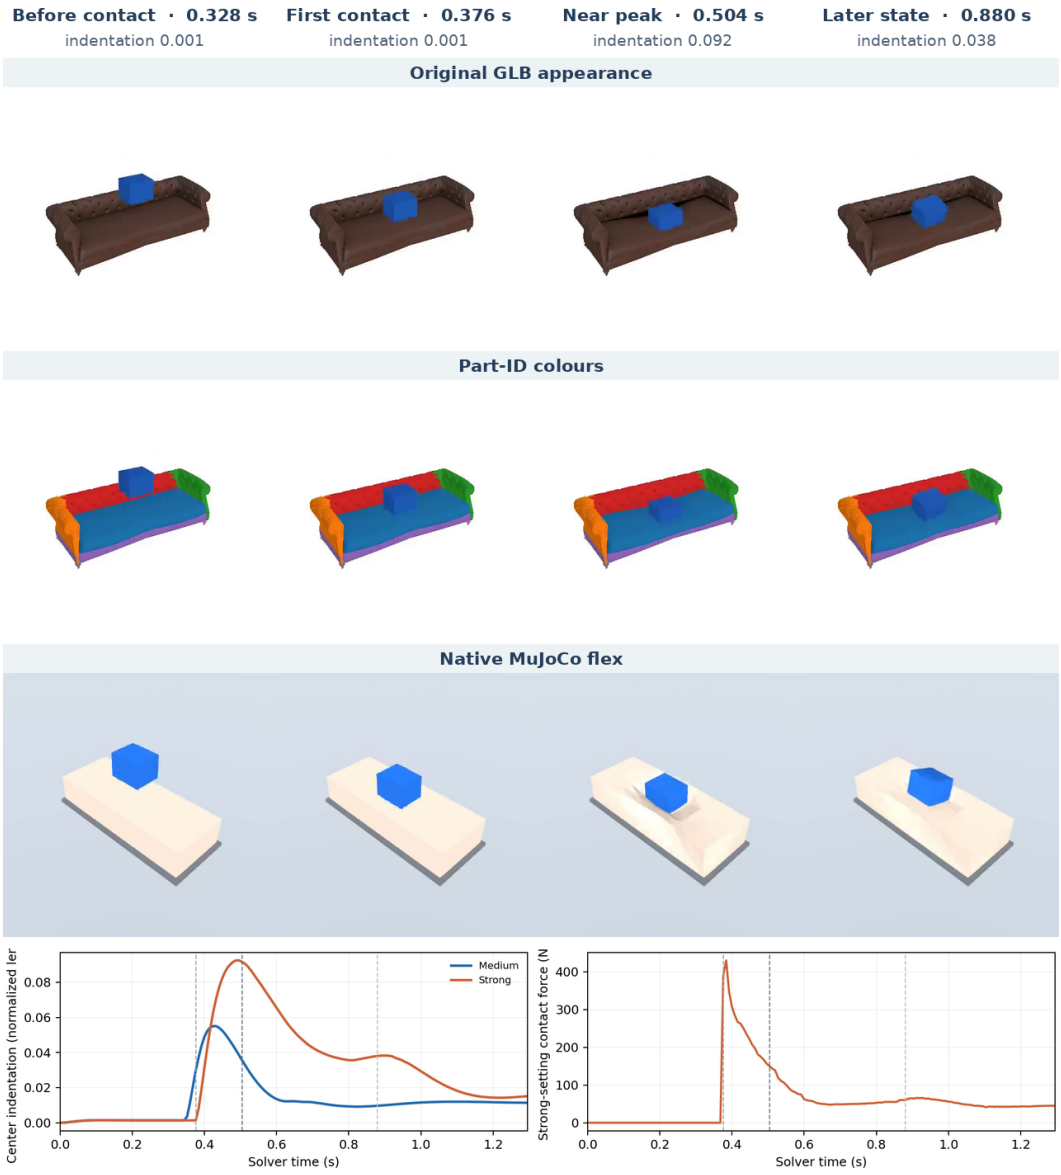

Figure 35: Saved rigid-block contact process for the held-out second sofa seat (model 58d54c95c3). Columns show the same saved-state indices before contact, at first recorded contact, nearest the solver's peak center-vertex indentation, and later in the response. Rows show original GLB appearance, distinct part-ID colours, and native MuJoCo block/flex geometry. Times are 0.328, 0.376, 0.504, 0.880 s, with center indentations 0.001, 0.001, 0.092, 0.038 normalized length units. Lower plots show center indentation under medium and strong authored loads and the strong-run normal contact force (peak 430.4 N). Medium and strong settings change both rigid mass and drop height, so this is a process comparison rather than an isolated stiffness test. The original GLB triangles are visually skinned from saved proxy motion; the native solver uses a coarse pinned volumetric flex, and predicted material priors do not set its parameters. All part-property predictions appear in Figure 19.

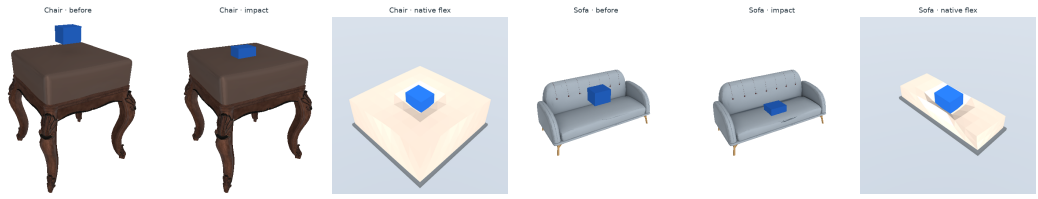

Figure 36: Compact qualitative comparison on original held-out chair and sofa GLBs before and during rigid-block contact, paired with the native MuJoCo flex state. Original source-appearance triangles are visually skinned from the saved solver motion. Parameters are authored; the panels do not validate predicted physical properties.

M.2 Simultaneous multi-body contact

Figure 37 uses the original bed from Figure 11 and three independent free rigid boxes, initially released at the same time above separate regions. Their identical box volume is evaluated at an illustrative scale of 0.8 m per normalized length unit. Particleboard, thermoplastic, and borosilicate source-bucket densities of 710, 1200, and 2230 kg/m^3 give rigid masses of 1.99, 3.36, and 6.25 kg. All three contact the same pinned tetrahedral flex body; their local center vertices reach 0.074, 0.131, and 0.183 normalized length units of downward displacement, with 529 contact steps recorded per box. The figure and video synchronize the original source-appearance GLB, distinct part-ID colours, and a native MuJoCo view. The native replay agrees with saved flex vertices to within 3×10^{-8} normalized units. Since the contact locations differ, the three local displacements alone do not identify a density effect. The source GLB triangles are visually skinned, while the flex proxy alone is the collision body. Source-bucket E and ν are displayed in Table 24, but do not enter MuJoCo's rigid-body solver; the bed flex parameters and scale assumption are authored. This demo does not validate the material predictor.



Figure 37: Three separate rigid blocks fall on the mattress/blanket region of one held-out PartNeXt bed. The first three panels show the original GLB triangles and its source appearance before and at peak contact; the right panel is a native MuJoCo rigid–flex render. Part-ID colours distinguish the bed components; blue, orange and purple distinguish the falling blocks.

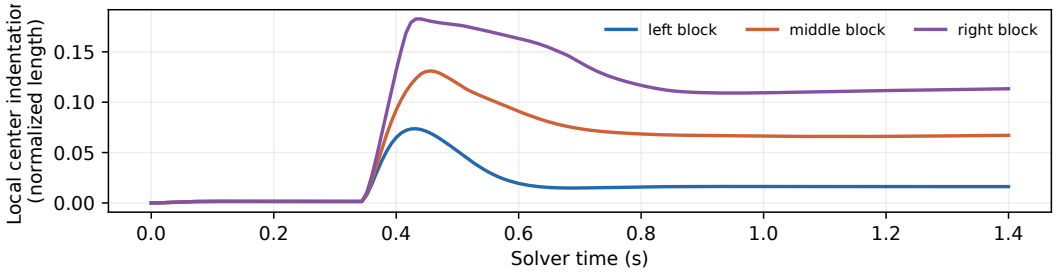

Figure 38: Saved MuJoCo flex-surface indentation under each of the three blocks in Figure 37. The curves are solver states in normalized model units, not measured bed displacements.

Same-location density control. To separate rigid mass from contact location, we rerun one box at the identical central impact point with the same box size, drop height, bed flex, and solver settings. Only source-bucket density, and hence rigid mass under the common assumed scale, changes; display colour is immaterial. The peak center-vertex indentations rise from 0.069 to 0.116 to 0.175 normalized units (Table 24). Figure 39 shows both the native MuJoCo peak states and the saved displacement curves. This establishes a controlled mass effect in the authored proxy simulator, not the accuracy of PartNeXt-Mat's material-property predictions or a physically calibrated bed. The offline supplement includes the synchronized native-state video and XML/state files for all three runs.

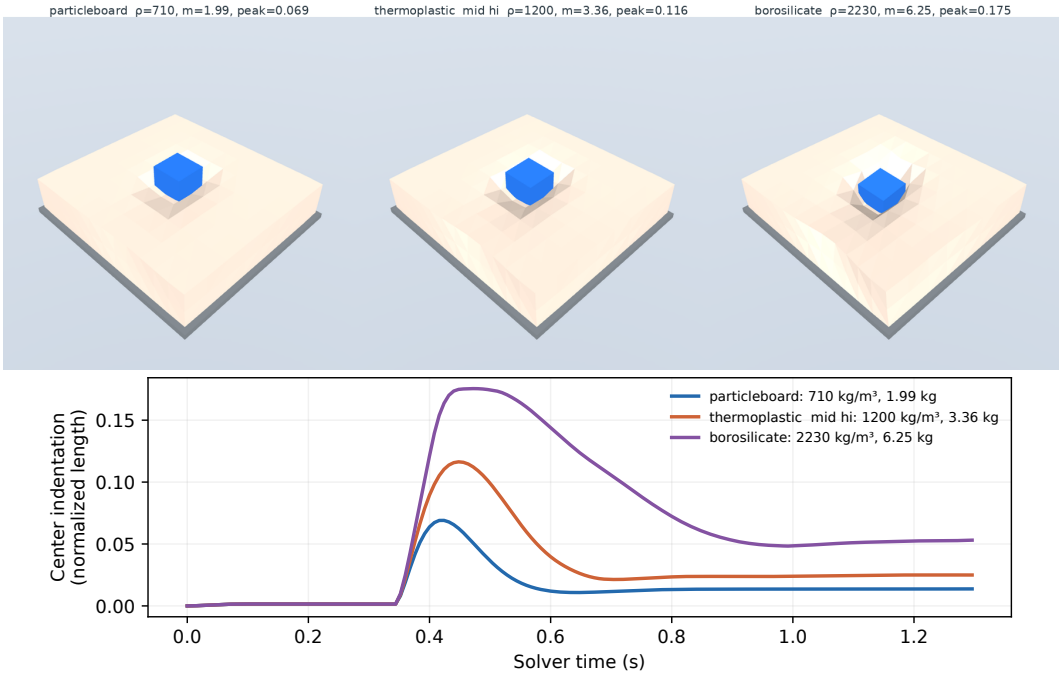

Figure 39: Controlled rigid-density sweep at one fixed bed-proxy contact point. Native MuJoCo peaks (top) and center-flex displacement histories (bottom) use identical geometry, drop, flex and solver parameters; only rigid mass changes according to the displayed source-bucket density. The original GLB is not the flex collision mesh.

Rigid source bucket	ρ kg/m^3	E GPa*	ν*	mass kg	peak disp.
Particleboard	710	3.20	0.3665	1.99	0.069
Thermoplastic	1200	2.25	0.3643	3.36	0.116
Borosilicate	2230	62.6	0.2280	6.25	0.175

Table 24: Source-bucket references for the three rigid blocks. Mass is ρ times their common box volume under the illustrative 0.8 m/unit scale. Peak flex displacement is in normalized units. *Rigid-block E and ν are listed for transparency but do not affect the MuJoCo rigid contact calculation. None of these block properties was inferred by the benchmark model.

Same-location controls on a chair and sofa. We repeat the one-factor mass design on the held-out chair cushion and sofa seat used in the block-contact case studies. For each target, all three runs reuse one box size, launch height, impact point, pinned flex, and MuJoCo settings. The same source-bucket densities produce masses of 1.99, 3.36, and 6.25 kg under the explicit illustrative scale. A direct XML comparison confirms that the three scenes for each target differ only in the rigid collider's mass attribute; the rest flex vertices, sampled times, and tracked center vertex also coincide. Chair peak center indentations are 0.049, 0.089, and 0.174 normalized units; sofa peaks are 0.054, 0.089, and 0.135. All six runs record 479 solver steps with rigid–flex contact, and native video replays agree with saved vertices to within 3.0×10^{-8} normalized units. Figures 40 and 41 show native

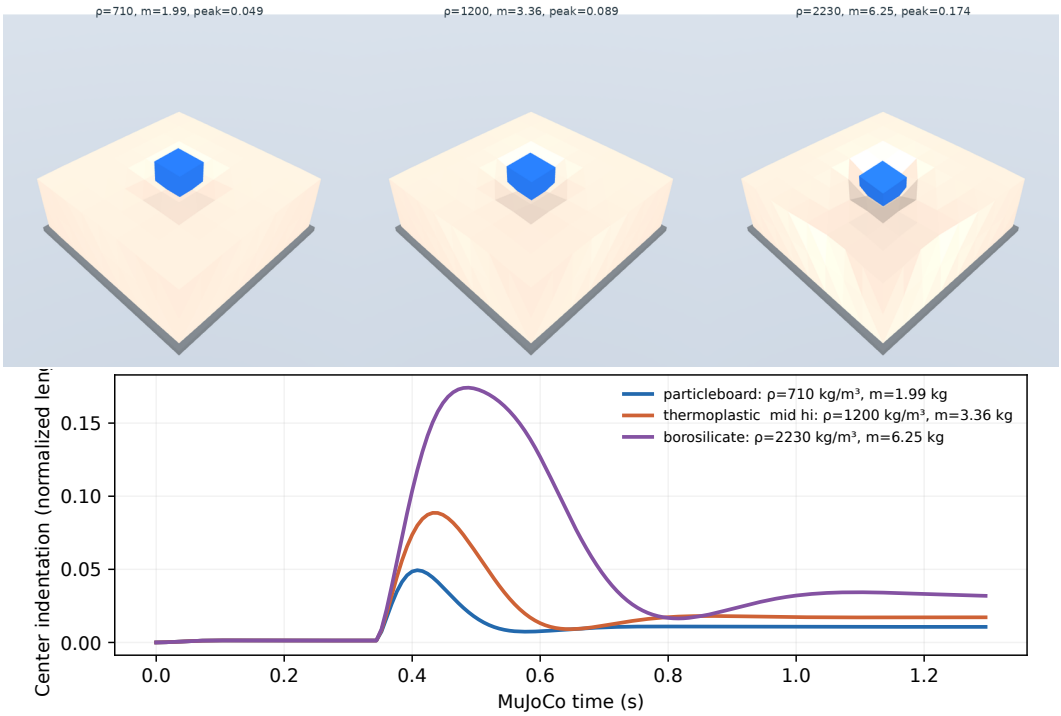

Figure 40: Same-location rigid-density sweep on the held-out chair cushion target (source model in Figure 13). Native MuJoCo peak states (top) and saved center-flex displacement histories (bottom) hold the box geometry, contact point, drop height, flex body, boundary conditions and solver settings fixed. Source-bucket densities 710, 1200, and 2230 kg/m^3 yield peak indentations 0.049, 0.089, 0.174 normalized model units. Only the rigid mass changes in MuJoCo; the associated source E and ν are metadata. The original GLB is visual context, not the collision mesh, and no predicted part property sets the flex stiffness. This is a controlled proxy response, not a material-prediction accuracy result.

peak shapes and complete displacement histories, while the web supplement includes synchronized videos, XML, states, and per-run provenance. The two source GLBs are original held-out PartNeXt assets, but the collision bodies are authored proxies; these controlled responses do not validate the benchmark model's material predictions.

Target	710 kg/m^3	1200 kg/m^3	2230 kg/m^3	High/low
Bed mattress	0.069	0.116	0.175	2.54
Chair cushion	0.049	0.089	0.174	3.53
Sofa seat	0.054	0.089	0.135	2.50

Table 25: Peak center-vertex indentation in normalized model units for three independent same-location density sweeps. Within each target, rigid mass is the only varied solver parameter; each target has its own authored flex and contact geometry. The ratios compare two loads within one proxy, not mechanics across real furniture or accuracy of predicted PartNeXt-Mat properties.

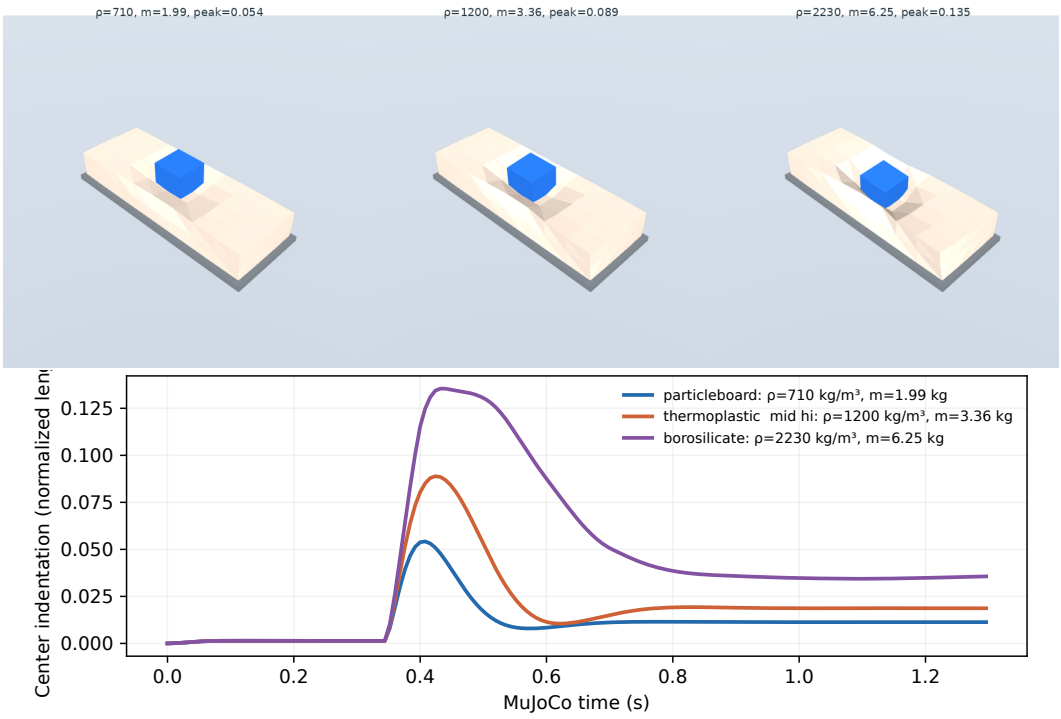

Figure 41: Same-location rigid-density sweep on the held-out sofa seat target (source model in Figure 20). Native MuJoCo peak states (top) and saved center-flex displacement histories (bottom) hold the box geometry, contact point, drop height, flex body, boundary conditions and solver settings fixed. Source-bucket densities 710, 1200, and 2230 kg/m^3 yield peak indentations 0.054, 0.089, 0.135 normalized model units. Only the rigid mass changes in MuJoCo; the associated source E and ν are metadata. The original GLB is visual context, not the collision mesh, and no predicted part property sets the flex stiffness. This is a controlled proxy response, not a material-prediction accuracy result.

N SEPARATELY SUPPLIED STRUCTURAL-ANCHOR BUNDLES

The companion package contains 88 simulation-ready bundles across 20 source categories: 66 have one FEM body, 14 have two, five have three, and three have four. Each includes a rigid assembly, one or more deformable surface meshes, material values, a constrained-surface mask, and source/anchor/rest-state renders. The offline supplement exposes all 88 source images, masks, rest renders, and source material tables. Their (ρ, E, ν) values are the package's source annotations, not predictions from the PartNeXt-Mat model; its object identifiers and source collection differ from the PartNeXt-Mat benchmark. We do not use these images as evidence of dynamic accuracy because the companion ZIP excludes the previously generated videos and vertex snapshots. An attempted local Genesis rerun rendered the rest state but did not complete a verified dynamic trajectory in the available environment. This section documents asset and boundary-condition diversity only.

O HELD-OUT PER-PART PROPERTY PREDICTIONS

The five cases below are selected by a fixed independent object hash: the first ID-test asset within each of Sofa, Umbrella, Chair, Lamp, and Handbag having 4–12 annotated parts and at least three released material classes. Selection uses neither predictions nor errors. Every one of their 43 parts appears in a separate table, including invisible parts and mistakes. The reported material is the class of the selected fine property prototype from the frozen fused Set Attention seed-zero model. This makes each applicable displayed (ρ, E, ν) internally consistent with its displayed material; the separately evaluated coarse-class prediction can differ. Fabric rows instead display the released packing-density and membrane-modulus template with a dagger and remain masked from volumetric evaluation. These are genuine predictions on held-out PartNeXt assets; the preceding real-photo plants and the separate structural-anchor bundles are not presented as PartNeXt-Mat test predictions.

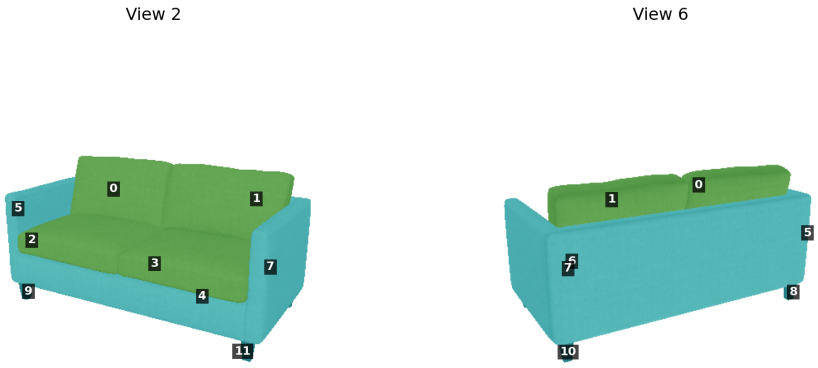

Figure 42: Held-out Sofa asset. Two fixed views of the original PartNeXt GLB show the frozen model's selected property-prototype class by part colour. Part IDs map to Table 26.

Table 26: Per-part predictions for the held-out Sofa in Figure 42. Values are released material priors, not measured asset mechanics.

ID	Part	Released	Predicted grade class	ρ (kg/m^3)	E (GPa)	ν
0	Backrest Cushion	fabric	fabric	$307.3^\dagger$	$0.00012^\dagger$	0.3612
1	Backrest Cushion	fabric	fabric	$307.3^\dagger$	$0.00012^\dagger$	0.3612
10	Leg	metal	wood	560	9.215	0.3665
11	Leg	metal	wood	560	9.215	0.3665
2	Seat Cushion	fabric	fabric	$307.3^\dagger$	$0.00012^\dagger$	0.3612
3	Seat Cushion	fabric	fabric	$307.3^\dagger$	$0.00012^\dagger$	0.3612
4	Seat Cushion Support Box	wood	wood	560	9.215	0.3665
5	Armrest Support Box	wood	wood	560	9.215	0.3665
6	Backrest Frame	wood	wood	560	9.215	0.3665
7	Armrest Support Box	wood	wood	560	9.215	0.3665
8	Leg	metal	wood	560	9.215	0.3665
9	Leg	metal	wood	560	9.215	0.3665

† Fabric numbers are the released packing-density and membrane-modulus template, not volumetric ρ, E targets; they are excluded from volumetric metrics and do not calibrate the illustrative flex proxies.

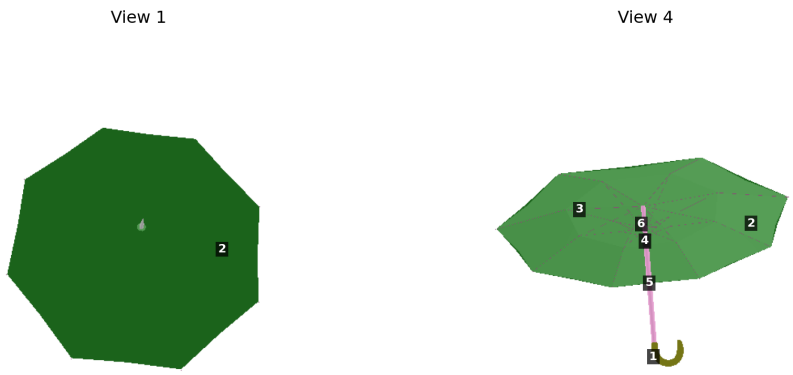

Figure 43: Held-out Umbrella asset. Two fixed views of the original PartNeXt GLB show the frozen model's selected property-prototype class by part colour. Part IDs map to Table 27.

Table 27: Per-part predictions for the held-out Umbrella in Figure 43. Values are released material priors, not measured asset mechanics.

ID	Part	Released	Predicted grade class	ρ (kg/m^3)	E (GPa)	ν
0	Tip	metal	plastic	1200	2.25	0.3643
1	Grip	rubber	rubber	1000	0.0088	0.49
2	Fabric	fabric	fabric	307.3†	0.00012†	0.3612
3	Stretcher	metal	metal	2710	71	0.33
4	Central Shaft	plastic	plastic	1200	2.25	0.3643
5	Central Rod	metal	metal	2710	71	0.33
6	Rib Rod	metal	metal	2710	71	0.33

† Fabric numbers are the released packing-density and membrane-modulus template, not volumetric ρ, E targets; they are excluded from volumetric metrics and do not calibrate the illustrative flex proxies.

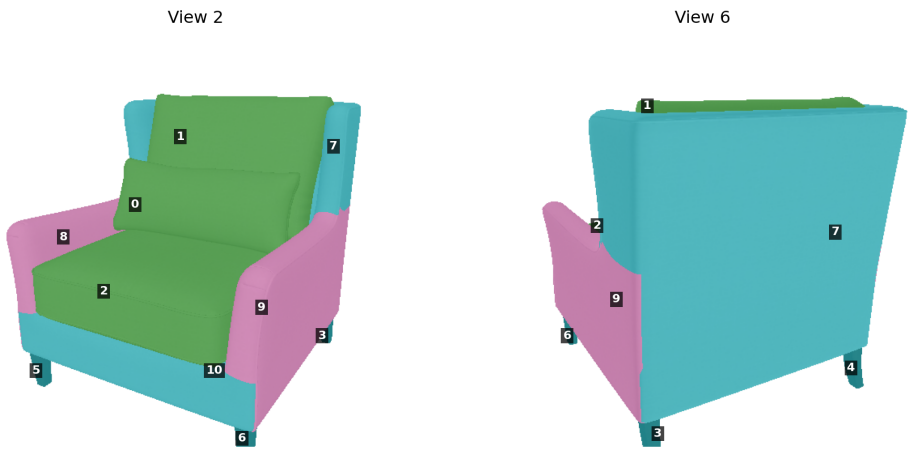

Figure 44: Held-out Chair asset. Two fixed views of the original PartNeXt GLB show the frozen model's selected property-prototype class by part colour. Part IDs map to Table 28.

Table 28: Per-part predictions for the held-out Chair in Figure 44. Values are released material priors, not measured asset mechanics.

ID	Part	Released	Predicted grade class	ρ (kg/m^3)	E (GPa)	ν
0	Cushion	fabric	fabric	307.3†	0.00012†	0.3612
1	Back Cushion	fabric	fabric	307.3†	0.00012†	0.3612
10	Frame	wood	wood	560	9.215	0.3665
2	Cushion	fabric	fabric	307.3†	0.00012†	0.3612
3	Leg	wood	wood	560	9.215	0.3665
4	Leg	wood	wood	560	9.215	0.3665
5	Leg	wood	wood	560	9.215	0.3665
6	Leg	wood	wood	560	9.215	0.3665
7	Integrated Frame	wood	wood	560	9.215	0.3665
8	Armrest Horizontal Rod	metal	metal	7850	203	0.29
9	Armrest Horizontal Rod	metal	metal	7850	203	0.29

† Fabric numbers are the released packing-density and membrane-modulus template, not volumetric ρ, E targets; they are excluded from volumetric metrics and do not calibrate the illustrative flex proxies.



Figure 45: Held-out Lamp asset. Two fixed views of the original PartNeXt GLB show the frozen model's selected property-prototype class by part colour. Part IDs map to Table 29.

Table 29: Per-part predictions for the held-out Lamp in Figure 45. Values are released material priors, not measured asset mechanics.

ID	Part	Released	Predicted grade class	ρ (kg/m^3)	E (GPa)	ν
0	Base	ceramic	metal	7850	203	0.29
1	Shade Cover	fabric	fabric	$307.3^{\dagger}$	$0.00012^{\dagger}$	0.3612
2	Connection Part	metal	metal	7850	196.5	0.29
3	Pole	metal	metal	7850	196.5	0.29

† Fabric numbers are the released packing-density and membrane-modulus template, not volumetric ρ, E targets; they are excluded from volumetric metrics and do not calibrate the illustrative flex proxies.



Figure 46: Held-out Handbag asset. Two fixed views of the original PartNeXt GLB show the frozen model's selected property-prototype class by part colour. Part IDs map to Table 30.

Table 30: Per-part predictions for the held-out Handbag in Figure 46. Values are released material priors, not measured asset mechanics.

ID	Part	Released	Predicted grade class	ρ (kg/m^3)	E (GPa)	ν
0	Main Handle	rubber	rubber	1000	0.0088	0.49
1	Main Handle	rubber	rubber	1000	0.0088	0.49
2	Bottom Stud	plastic	plastic	1048	0.6118	0.3643
3	Bottom Stud	plastic	plastic	1048	0.6118	0.3643
4	Bottom Stud	plastic	plastic	1048	0.6118	0.3643
5	Bottom Stud	plastic	plastic	1048	0.6118	0.3643
6	Main Part	metal	metal	8600	110	0.33
7	Main Handle	rubber	rubber	1000	0.0088	0.49
8	Main Handle	rubber	rubber	1000	0.0088	0.49

† Fabric numbers are the released packing-density and membrane-modulus template, not volumetric ρ, E targets; they are excluded from volumetric metrics and do not calibrate the illustrative flex proxies.

P BROADER HELD-OUT ASSET GALLERY

To expose more than a few favourable examples, we also render 88 held-out PartNeXt-Mat assets spanning 47 categories. Selection is fixed before inspecting predictions: for each ID-test category with a 4–12-part eligible object, we take the first object under an independent hash; a second object is taken from 41 of the 42 categories with two eligible objects, again chosen by an independent category hash. The gallery contains all 606 annotated parts in these objects. Each card colours the actual GLB triangles by the fine property prototype selected by the frozen fused Set Attention seed-zero model, while the following tables list every part's predicted material class and applicable (ρ, E, ν) values. Incorrect assignments remain visible. The 88 assets here are selected from the PartNeXt-Mat test split and are distinct from the 88 FEM bundles in Section N. No simulator behavior is inferred from these static gallery renders.

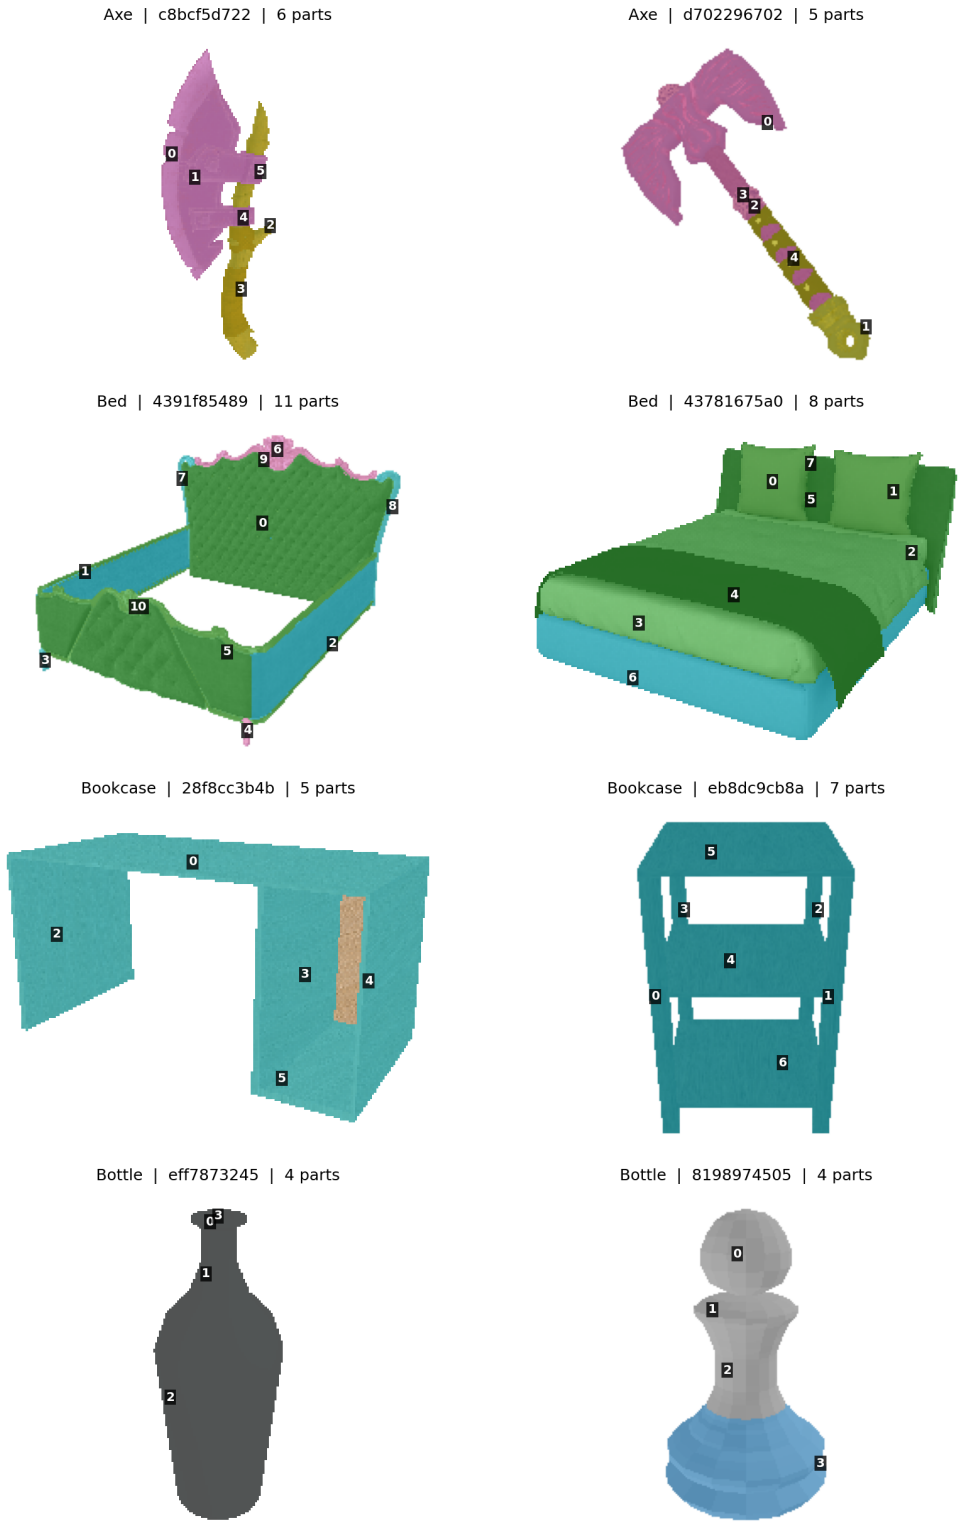

Figure 47: Held-out PartNeXt-Mat assets 1–8 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.

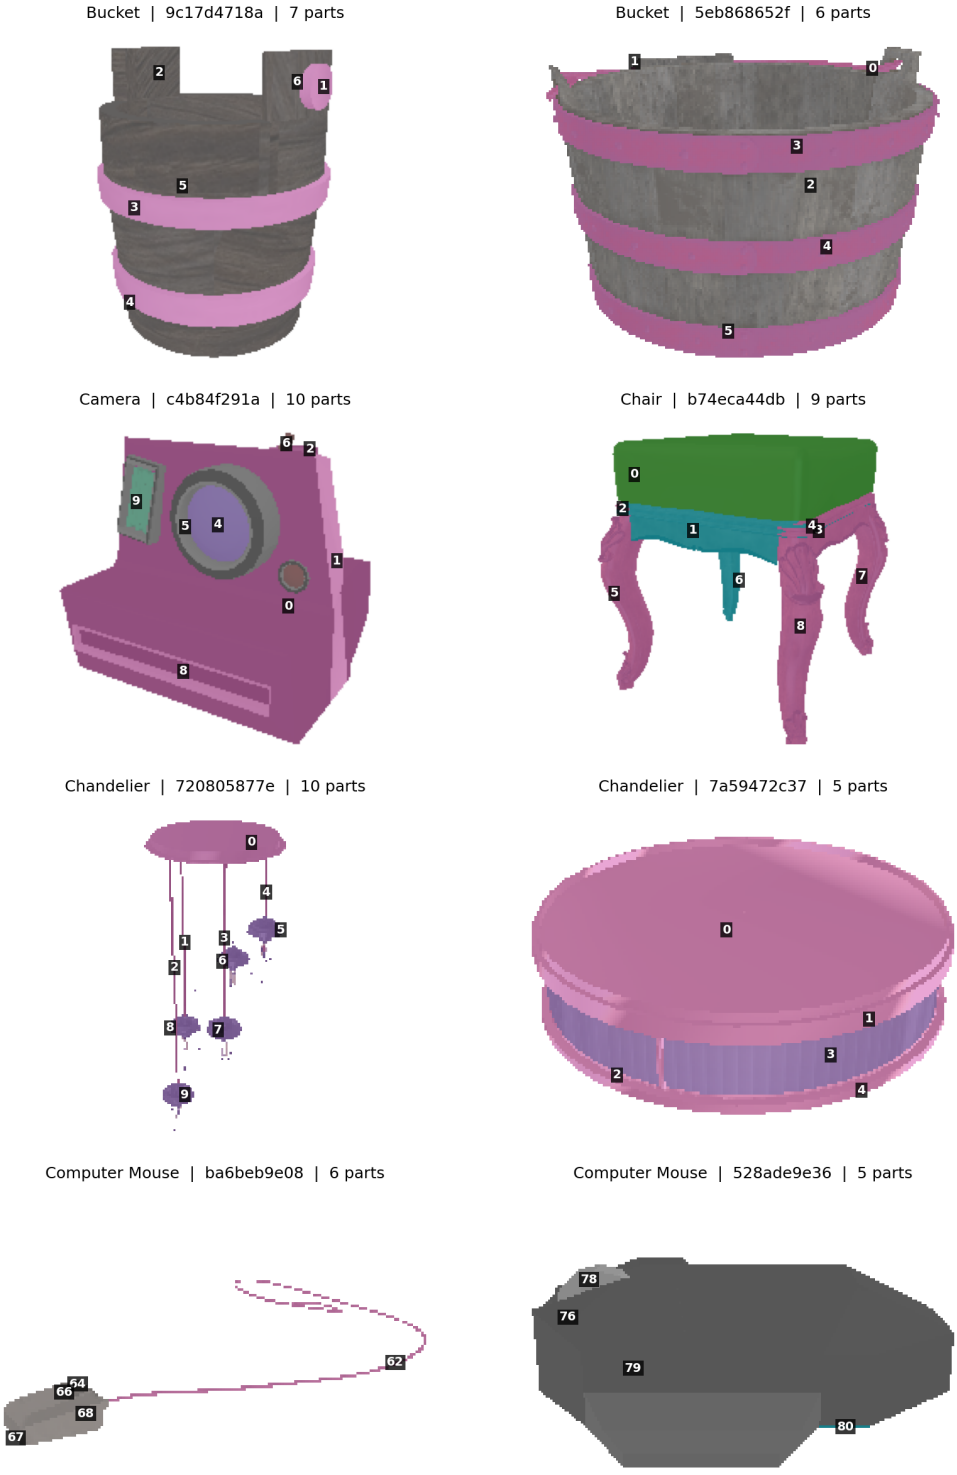

Figure 48: Held-out PartNeXt-Mat assets 9–16 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.



Figure 49: Held-out PartNeXt-Mat assets 17–24 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.

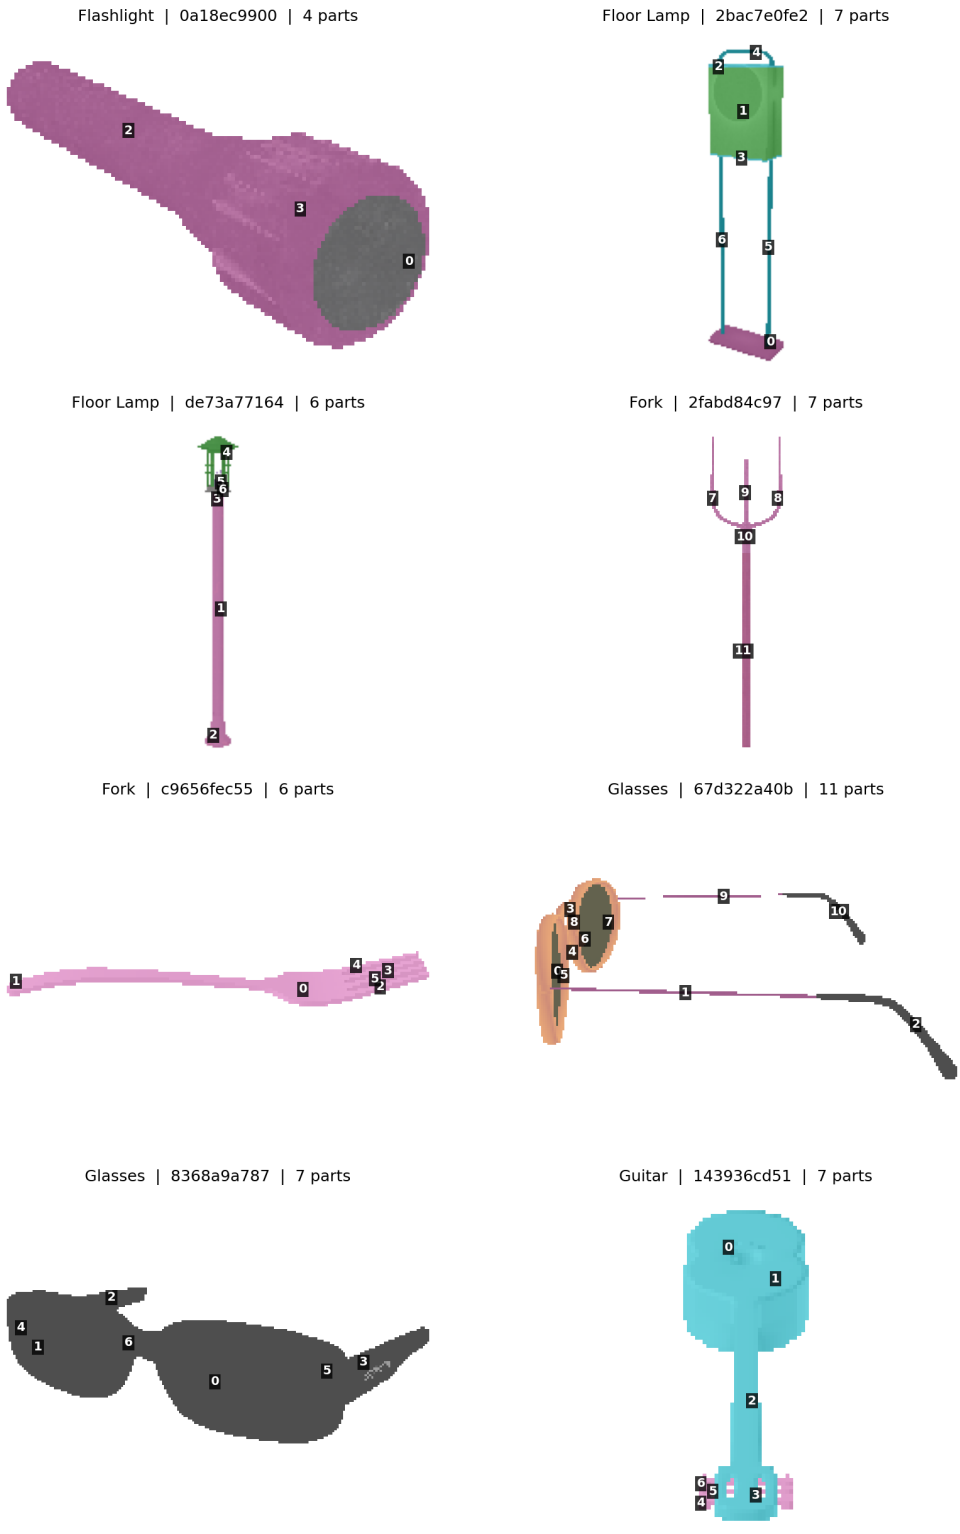

Figure 50: Held-out PartNeXt-Mat assets 25–32 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.



Figure 51: Held-out PartNeXt-Mat assets 33–40 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.

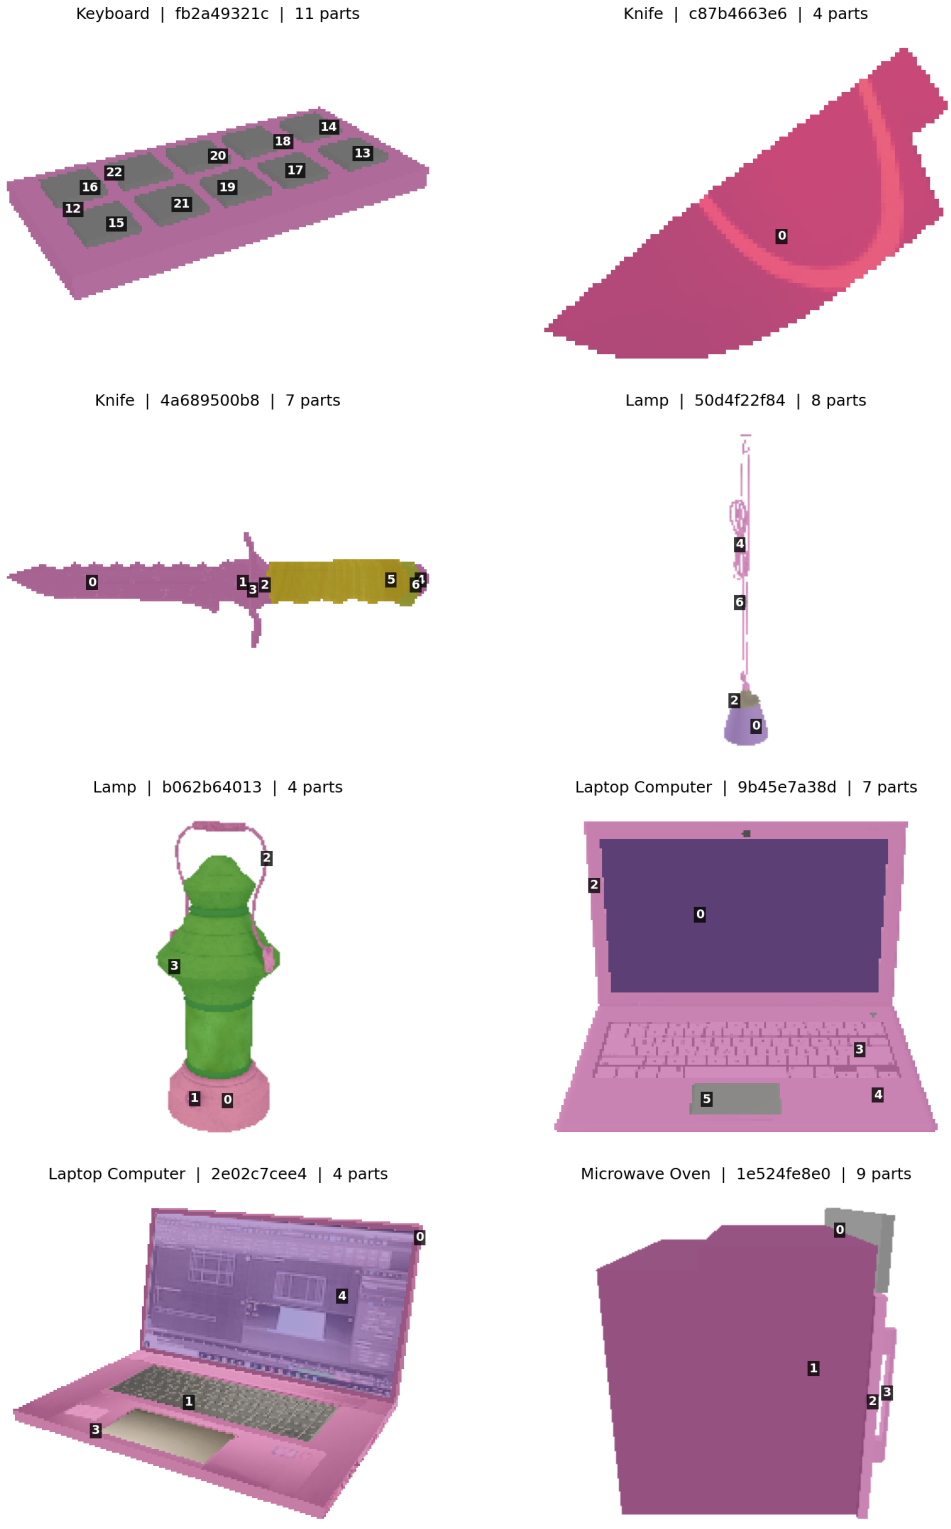

Figure 52: Held-out PartNeXt-Mat assets 41–48 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.

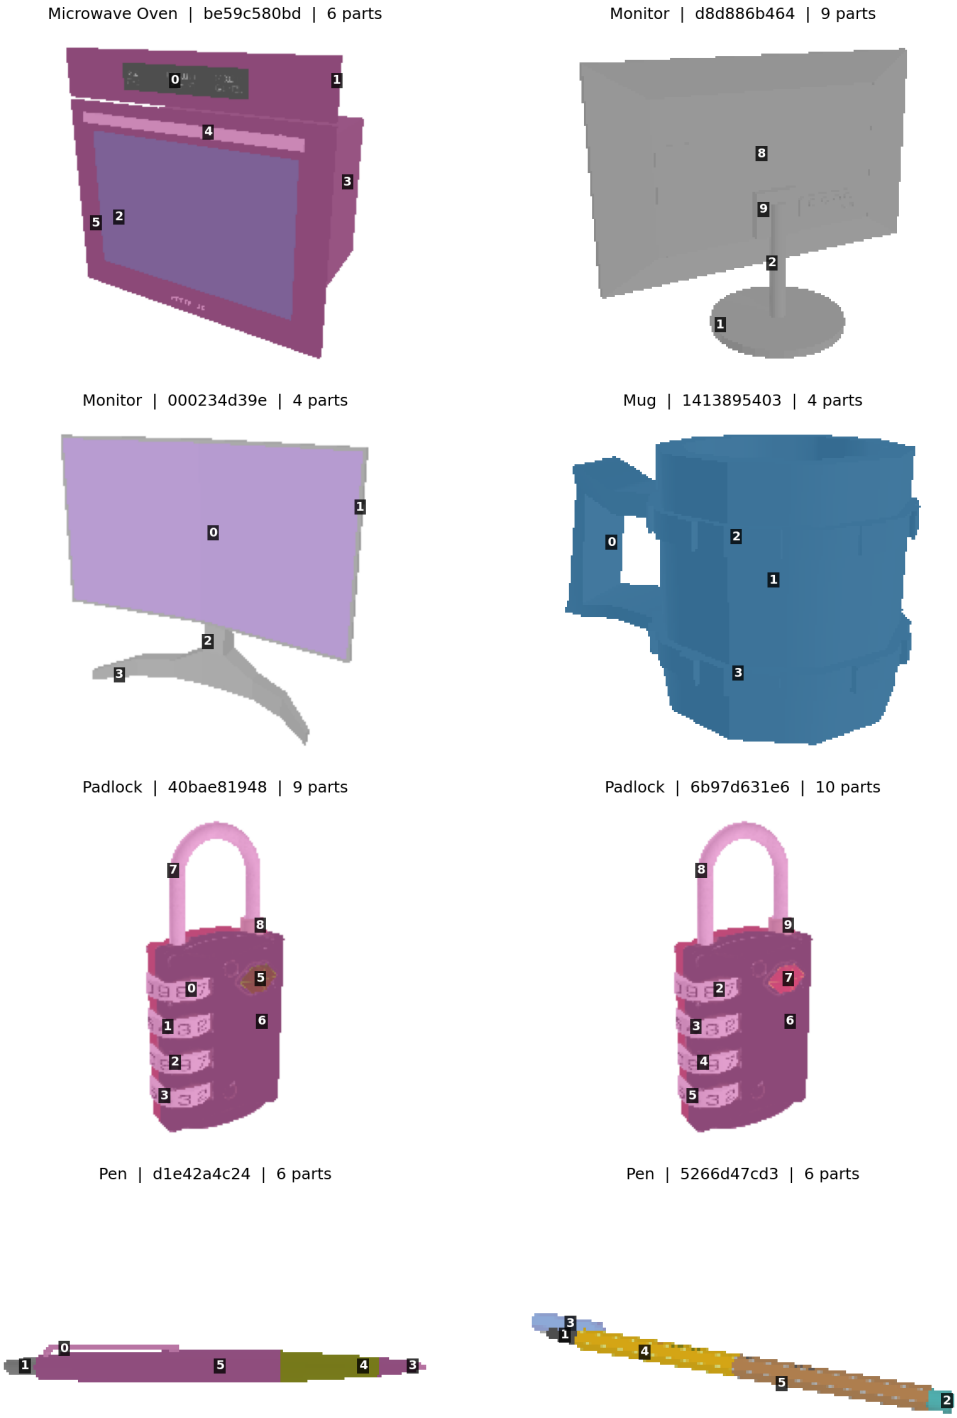

Figure 53: Held-out PartNeXt-Mat assets 49–56 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.

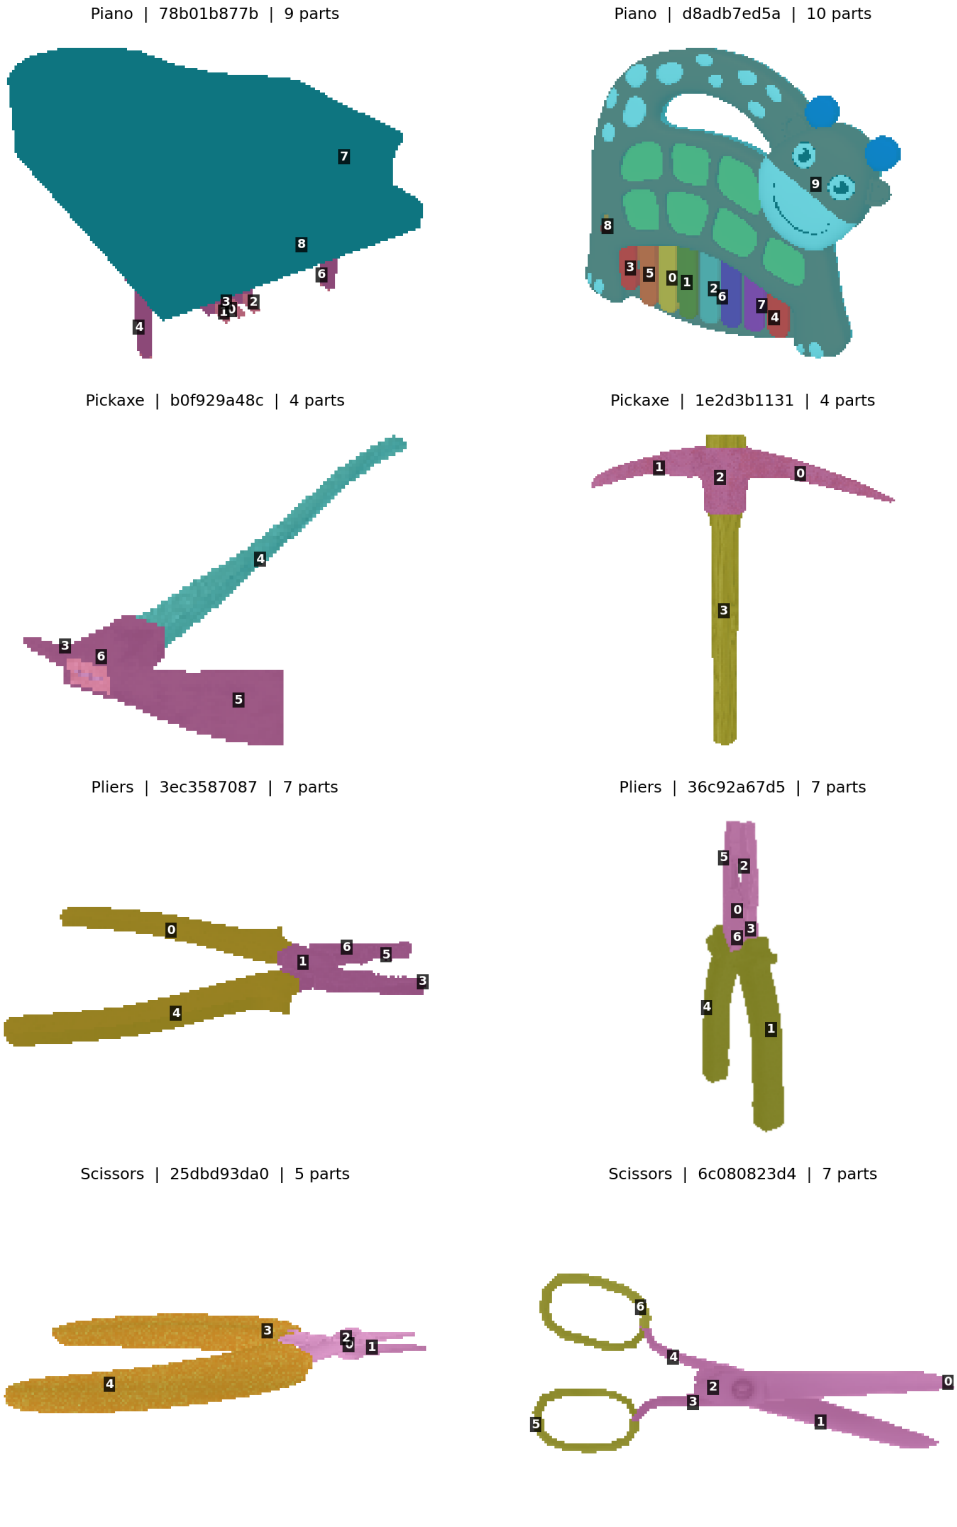

Figure 54: Held-out PartNeXt-Mat assets 57–64 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.



Figure 55: Held-out PartNeXt-Mat assets 65–72 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.



Figure 56: Held-out PartNeXt-Mat assets 73–80 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.

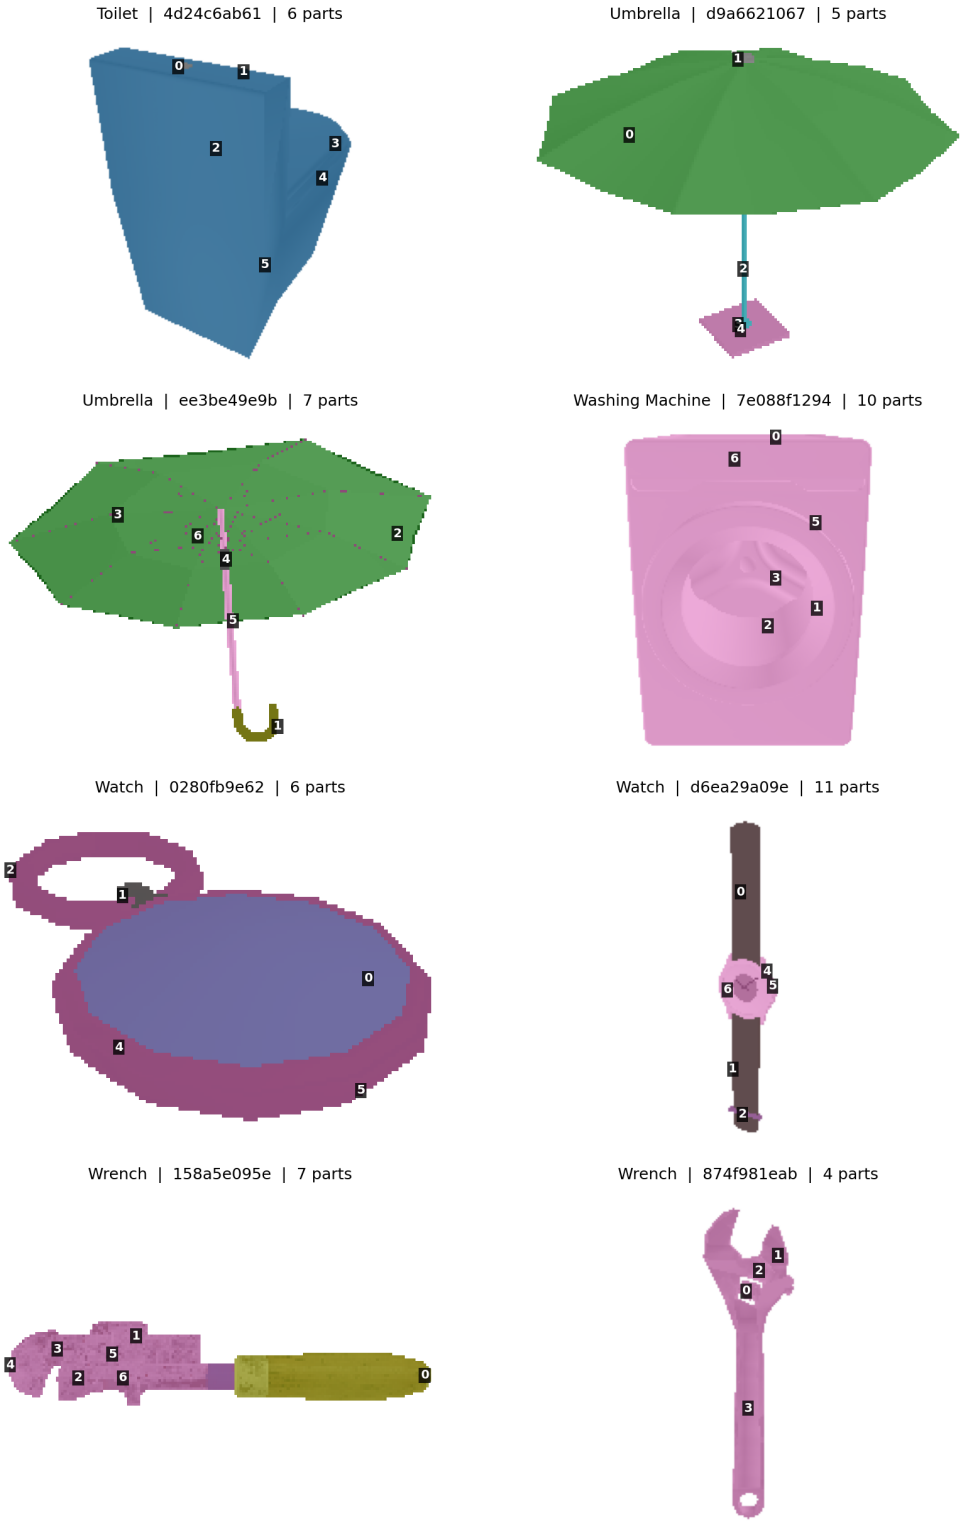

Figure 57: Held-out PartNeXt-Mat assets 81–88 of 88. Each card shows one fixed rendered GLB view, with visible part IDs and tint from the frozen model's selected material-property prototype. Per-part predictions and marked reference-only values follow in separate tables. Selection ignores predictions and errors. These are visualizations of PartNeXt assets, not FEM simulations.

Q COMPLETE PROPERTIES FOR THE 88 DISPLAYED ASSETS

For each part, the predicted class and applicable triple below come from the same selected fine prototype. A dagger marks the released fabric template entries: the reported density follows a bulk-packing convention and the reported modulus a membrane convention. These numbers are displayed for transparency, but are masked as volumetric ρ and E targets and must not be passed directly to a volumetric FEM/MPM solver. A double dagger marks 24 parts whose selected ceramic prototype is label-only. Their numeric cells show the separate whiteware corpus reference (2350 kg/m^3, 75.025 GPa, 0.204) solely to make the class prior visible; it is neither assigned to those individual assets nor a numerical model prediction. The fixed 88-asset selection is unchanged. Unmarked values are released-prior predictions, not measurements of individual objects. The JSON supplement retains the actual missing numerical outputs, coarse-class head, prototype identifier, and per-quantity applicability mask for every row.

1. Axe: c8bcf5d722bf4c6ca6e1f1f4383d3cb7 (6 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Blade	metal	metal	7850	203	0.29
1	Connector	metal	metal	7850	203	0.29
2	Grip	rubber	rubber	1000	0.0088	0.49
3	Anti-Slip Material	rubber	rubber	1000	0.0088	0.49
4	Connector	metal	metal	7850	196.5	0.29
5	Connector	metal	metal	7850	203	0.29

2. Axe: d702296702684cd4b52008adccf09011 (5 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Blade	metal	metal	7850	203	0.29
1	Handle End	rubber	rubber	1000	0.0088	0.49
2	Shaft	composite	metal	8600	110	0.33
3	Decoration	metal	metal	8600	110	0.33
4	Anti-Slip Material	rubber	rubber	1000	0.0088	0.49

3. Bed: 4391f85489ed41a5aa3dc77447b63b9b (11 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Soft Surface	fabric	fabric	307.3†	0.00012†	0.3612
1	Side Panel	wood	wood	660	5.135	0.3665
10	Other	fabric	fabric	307.3†	0.00012†	0.3612
2	Side Panel	wood	wood	660	5.135	0.3665
3	Leg	wood	wood	560	9.215	0.3665
4	Leg	wood	metal	8600	110	0.33
5	Soft Surface	fabric	fabric	307.3†	0.00012†	0.3612
6	Other	wood	metal	8600	110	0.33
7	Vertical Frame Bar	wood	wood	500	7.58	0.3665
8	Vertical Frame Bar	wood	wood	500	7.58	0.3665
9	Other	wood	metal	8600	110	0.33

4. Bed: 43781675a00c427aa7aa255f4225e69c (8 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Pillow	fabric	fabric	307.3†	0.00012†	0.3612
1	Pillow	fabric	fabric	307.3†	0.00012†	0.3612
2	Blanket	fabric	fabric	307.3†	0.00012†	0.3612
3	Mattress	fabric	fabric	307.3†	0.00012†	0.3612
4	Bed Scarf	fabric	fabric	307.3†	0.00012†	0.3612
5	Soft Surface	fabric	fabric	307.3†	0.00012†	0.3612
6	Bottom Panel	wood	wood	560	9.215	0.3665
7	Other	fabric	fabric	307.3†	0.00012†	0.3612

5. Bookcase: 28f8cc3b4b5141fca351835eae015946 (5 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Top Panel	wood	wood	560	9.215	0.3665
2	Side Panel	wood	wood	560	9.215	0.3665
3	Side Panel	wood	wood	560	9.215	0.3665
4	Side Panel	wood	wood	560	9.215	0.3665
5	Bottom Panel	wood	wood	795	3.56	0.3665

6. Bookcase: eb8dc9cb8aaf4c68b7b573eec4aa682f (7 parts, gallery view 1)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Post	metal	wood	710	3.2	0.3665
1	Post	metal	wood	710	3.2	0.3665
2	Post	metal	wood	710	3.2	0.3665
3	Post	metal	wood	710	3.2	0.3665
4	Shelf Body	wood	wood	710	3.2	0.3665
5	Top Panel	wood	wood	710	3.2	0.3665
6	Bottom Panel	metal	wood	560	9.215	0.3665

7. Bottle: eff7873245084790b5cb1e2ce1daf8b0 (4 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Mouth	plastic	plastic	1200	2.25	0.3643
1	Neck	plastic	plastic	1200	2.25	0.3643
2	Body	plastic	plastic	1200	2.25	0.3643
3	Other	plastic	plastic	1200	2.25	0.3643

8. Bottle: 81989745059d43499a903e4ab91179ac (4 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Cap Body	plastic	plastic	1200	2.25	0.3643
1	Mouth	plastic	plastic	1200	2.25	0.3643
2	Neck	plastic	plastic	1200	2.25	0.3643
3	Body	plastic	ceramic	2350	75.02	0.204

9. Bucket: 9c17d4718abf46f598de44ab26475bb9 (7 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Fixed Rivet	metal	metal	7850	203	0.29
1	Fixed Rivet	metal	metal	7850	203	0.29
2	Connection Main Body	plastic	plastic	1200	2.25	0.3643
3	Hoop	metal	metal	2710	71	0.33
4	Hoop	metal	metal	2710	71	0.33
5	Integrated Main Body	plastic	plastic	1200	2.25	0.3643
6	Connection Main Body	plastic	plastic	1200	2.25	0.3643

10. Bucket: 5eb868652fdc4531990e9e5454e00d89 (6 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Handle Main Body	metal	metal	8600	110	0.33
1	Handle Grip	rubber	plastic	1200	2.25	0.3643
2	Integrated Main Body	metal	plastic	1200	2.25	0.3643
3	Hoop	metal	metal	8600	110	0.33
4	Hoop	metal	metal	2710	71	0.33
5	Hoop	metal	metal	8600	110	0.33

11. Camera: c4b84f291a12497b833ad222a8159c27 (10 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Base	metal	metal	2710	71	0.33
1	Casing	plastic	metal	2710	71	0.33
2	Casing	plastic	metal	2710	71	0.33
3	Lens Element	plastic	glass	2250	65.7	0.228
4	Lens Element	glass	glass	2250	65.7	0.228
5	Casing	plastic	plastic	1200	2.25	0.3643
6	Button	plastic	plastic	1200	2.25	0.3643
7	Button	plastic	plastic	1200	2.25	0.3643
8	Other	plastic	metal	2710	71	0.33
9	Screen Body	plastic	plastic	1200	2.25	0.3643

12. Chair: b74eca44dbd84f28a11aff514f86996c (9 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Cushion	fabric	fabric	$307.3^\dagger$	$0.00012^\dagger$	0.3612
1	Frame	wood	wood	560	9.215	0.3665
2	Frame	wood	metal	8600	127.6	0.27
3	Frame	wood	wood	560	9.215	0.3665
4	Frame	wood	metal	8600	127.6	0.27
5	Leg	wood	metal	8600	127.6	0.27
6	Leg	wood	wood	560	9.215	0.3665
7	Leg	wood	metal	8600	127.6	0.27
8	Leg	wood	metal	8600	127.6	0.27

13. Chandelier: 720805877e7a42b4856283df01f4af9e (10 parts, gallery view 5)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Base	metal	metal	7850	203	0.29
1	Chain	metal	metal	7850	203	0.29
2	Chain	metal	metal	7850	203	0.29
3	Chain	metal	metal	7850	203	0.29
4	Chain	metal	metal	7850	203	0.29
5	Shade Cover	fabric	glass	2230	62.6	0.228
6	Shade Cover	fabric	glass	2230	62.6	0.228
7	Shade Cover	fabric	glass	2230	62.6	0.228
8	Shade Cover	fabric	glass	2230	62.6	0.228
9	Shade Cover	fabric	glass	2230	62.6	0.228

14. Chandelier: 7a59472c374642afab07e621fad82066 (5 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Base	metal	metal	8600	110	0.33
1	Support Ring	metal	metal	8600	110	0.33
2	Support Ring	metal	metal	8600	110	0.33
3	Shade Cover	fabric	glass	2230	62.6	0.228
4	Lamp Holder	ceramic	metal	7850	196.5	0.29

15. Computer Mouse: ba6beb9e08e44855bd16d5d040c6536d (6 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
62	Wire	metal	metal	7850	203	0.29
64	Left Button	plastic	plastic	1200	2.25	0.3643
65	Connector	metal	plastic	1200	2.25	0.3643
66	Right Button	plastic	plastic	1200	2.25	0.3643
67	Bottom Frame	metal	plastic	1200	2.25	0.3643
68	Upper Shell Body	plastic	plastic	1200	2.25	0.3643

16. Computer Mouse: 528ade9e36cd4906a764f2db0515356c (5 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
76	Left Button	metal	plastic	1200	2.25	0.3643
77	Right Button	metal	plastic	1200	2.25	0.3643
78	Scroll Wheel	plastic	plastic	1200	2.25	0.3643
79	Upper Shell Body	metal	plastic	1200	2.25	0.3643
80	Bottom Frame	metal	wood	560	9.215	0.3665

17. Controller: ae34958d91ef4689a8545fe8d3b22f9f (8 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Buttons	plastic	plastic	1200	2.25	0.3643
1	Buttons	plastic	plastic	1200	2.25	0.3643
2	Buttons	plastic	plastic	1200	2.25	0.3643
3	Buttons	plastic	plastic	1200	2.25	0.3643
4	Buttons	plastic	plastic	1200	2.25	0.3643
5	Buttons	plastic	plastic	1200	2.25	0.3643
6	Directional Pad	plastic	plastic	1200	2.25	0.3643
7	Casing	plastic	plastic	1200	2.25	0.3643

18. Controller: cff2d3483c4a452aaa9201b2a1b7d984 (6 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Buttons	plastic	plastic	1200	2.25	0.3643
1	Buttons	plastic	plastic	1200	2.25	0.3643
2	Thumb Cap	plastic	plastic	1200	2.25	0.3643
3	Base	plastic	plastic	1200	2.25	0.3643
4	Stick	plastic	plastic	1200	2.25	0.3643
5	Casing	plastic	plastic	1200	2.25	0.3643

19. Cup: f9401a8c4f7d4f738d2da663ce6c67f2 (5 parts, gallery view 1)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Bottom	glass	metal	7850	203	0.29
1	Foot	glass	metal	7850	203	0.29
2	Rim	glass	metal	7850	196.5	0.29
3	Main Body	ceramic	metal	7850	196.5	0.29
4	Main Body	ceramic	metal	7850	196.5	0.29

20. Cup: 988f9944239c4b04986e33632b13804d (6 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Handle	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
1	Handle	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
2	Bottom	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
3	Foot	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
4	Rim	ceramic	metal	7850	196.5	0.29
5	Main Body	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$

21. Door: 6467eae771d94f4e8224f57f83a2a8c3 (9 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Door Frame	metal	metal	8600	110	0.33
1	Other	metal	metal	7800	196.5	0.29
2	Panel	wood	wood	795	3.56	0.3665
3	Keyhole	metal	metal	7800	196.5	0.29
4	Keyhole	metal	metal	7800	196.5	0.29
5	Handle Body	plastic	metal	7800	196.5	0.29
6	Connecting Rod	metal	metal	7850	203	0.29
7	Handle Body	plastic	metal	7800	196.5	0.29
8	Connecting Rod	metal	metal	7850	203	0.29

22. Fan: ff5b074af2c44084b1bfa882ac068197 (4 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Column	metal	metal	7850	203	0.29
1	Base Plate	metal	metal	7850	203	0.29
2	Bracket Body	metal	metal	7850	203	0.29
3	Casing	metal	metal	7850	203	0.29

23. Fan: 25f6e6a9b791458da09e75dc60e47a6f (7 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Impeller Cover	metal	metal	8600	110	0.33
1	Anti-Vibration Feet	metal	metal	8600	127.6	0.27
2	Blade Body	metal	plastic	1200	2.25	0.3643
3	Blade Body	metal	plastic	1200	2.25	0.3643
4	Blade Body	metal	plastic	1200	2.25	0.3643
5	Blade Connection	metal	metal	7850	203	0.29
6	Casing Body	metal	metal	7850	196.5	0.29

24. Flashlight: 467f78253a974f7c98506a85db5eebb6 (9 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Lens	glass	glass	2250	65.7	0.228
1	Bulb	glass	glass	2230	62.6	0.228
2	Bezel	metal	metal	2710	71	0.33
3	Switch Area	metal	metal	2710	71	0.33
4	Button	metal	plastic	1200	2.25	0.3643
5	Button	plastic	metal	8600	110	0.33
6	Connection Part	metal	metal	2710	71	0.33
7	Tail Cap	metal	metal	2710	71	0.33
8	Body	metal	metal	2710	71	0.33

25. Flashlight: 0a18ec9900b94d3ab1efe442af57c0a3 (4 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Lens	glass	plastic	1200	2.25	0.3643
1	Tail Cap	metal	plastic	1200	2.25	0.3643
2	Body	metal	metal	2710	71	0.33
3	Connection Part	metal	metal	2710	71	0.33

26. Floor Lamp: 2bac7e0fe2d54d9d8d8ea9d01d1a6310 (7 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Base Body	metal	metal	7850	203	0.29
1	Shade Cover	fabric	fabric	$307.3^\dagger$	$0.00012^\dagger$	0.3612
2	Shade Connector Rod	wood	wood	560	9.215	0.3665
3	Shade Connector Rod	wood	wood	560	9.215	0.3665
4	Support Pole	metal	wood	560	9.215	0.3665
5	Leg	wood	wood	560	9.215	0.3665
6	Leg	wood	wood	560	9.215	0.3665

27. Floor Lamp: de73a77164884a6a95dd54201e328962 (6 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
1	Pole	wood	metal	7850	196.5	0.29
2	Base Body	fabric	metal	7850	203	0.29
3	Connector	plastic	metal	7850	196.5	0.29
4	Shade Cover	fabric	fabric	$307.3^\dagger$	$0.00012^\dagger$	0.3612
5	Bulb	glass	glass	2230	62.6	0.228
6	Lamp Socket	plastic	plastic	1200	2.25	0.3643

28. Fork: 2fabd84c97ff40e6b68ca7a1f77bb3f9 (7 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
10	Connector	metal	metal	7800	196.5	0.29
11	Shaft	wood	metal	8600	110	0.33
12	Connector	metal	metal	8600	110	0.33
13	Connector	metal	metal	8600	110	0.33
7	Tines	metal	metal	7850	203	0.29
8	Tines	metal	metal	7850	203	0.29
9	Tines	metal	metal	7850	203	0.29

29. Fork: c9656fec559348c396fae8670927667e (6 parts, gallery view 1)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Neck	metal	metal	7850	203	0.29
1	Shaft	metal	metal	7850	196.5	0.29
2	Tines	metal	metal	7850	203	0.29
3	Tines	metal	metal	7850	203	0.29
4	Tines	metal	metal	7850	203	0.29
5	Tines	metal	metal	7850	203	0.29

30. Glasses: 67d322a40b9a494fab372d4f0ed0dff2 (11 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Lens	plastic	plastic	1200	2.25	0.3643
1	Temple Arm	plastic	metal	7800	196.5	0.29
10	Temple Tip	plastic	plastic	1200	2.25	0.3643
2	Temple Tip	plastic	plastic	1200	2.25	0.3643
3	Nose Bridge	metal	metal	7850	203	0.29
4	Nose Support	plastic	metal	7850	203	0.29
5	Frame	metal	metal	7850	203	0.29
6	Nose Support	plastic	metal	7850	203	0.29
7	Lens	plastic	plastic	1200	2.25	0.3643
8	Frame	metal	metal	7850	203	0.29
9	Temple Arm	plastic	metal	7850	196.5	0.29

31. Glasses: 8368a9a7875641e9a3c9c6b7cf32b73b (7 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Lens	plastic	plastic	1200	2.25	0.3643
1	Lens	plastic	plastic	1200	2.25	0.3643
2	Temple Arm	plastic	plastic	1200	2.25	0.3643
3	Temple Arm	plastic	plastic	1200	2.25	0.3643
4	Frame	plastic	plastic	1200	2.25	0.3643
5	Frame	plastic	plastic	1200	2.25	0.3643
6	Nose Bridge	plastic	plastic	1200	2.25	0.3643

32. Guitar: 143936cd519948788ec704dd44ede68a (7 parts, gallery view 1)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Bridge Plate	metal	wood	560	9.215	0.3665
1	Body Frame	wood	wood	560	9.215	0.3665
2	Neck Body	wood	wood	560	9.215	0.3665
3	Headstock Body	wood	wood	560	9.215	0.3665
4	Nut	metal	metal	7850	203	0.29
5	Nut	metal	metal	7850	203	0.29
6	Nut	metal	metal	7850	203	0.29

33. Guitar: e62d47f2067c4ded96c16e4695113633 (7 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Strings	metal	metal	7850	203	0.29
1	Strings	metal	metal	7850	203	0.29
2	Strings	metal	metal	8600	110	0.33
3	Strings	metal	metal	7850	203	0.29
4	Body Frame	wood	metal	8600	127.6	0.27
5	Headstock Body	wood	metal	7850	203	0.29
6	Neck Body	wood	metal	8600	127.6	0.27

34. Hammer: 03b36c5f75eb46999ca2269b1d3321a5 (4 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Shaft	metal	rubber	1000	0.0088	0.49
1	Hammer Head	metal	metal	7850	203	0.29
2	Hammer Head	metal	metal	7850	203	0.29
3	Connector	metal	metal	7850	203	0.29

35. Hammer: 9ebd760261d5427dab8c5c546af3dba0 (6 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Hammer Head	metal	metal	7850	203	0.29
1	Connector	metal	metal	8600	127.6	0.27
2	Hammer Head	metal	metal	7850	203	0.29
3	Shaft	wood	metal	8600	110	0.33
4	Decoration	metal	rubber	1000	0.0088	0.49
5	Handle End	rubber	rubber	1000	0.0088	0.49

36. Handbag: 0447c8ac137b4b77a4873d0e2214cb0f (8 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Pocket	fabric	fabric	307.3†	0.00012†	0.3612
1	Pocket	fabric	fabric	307.3†	0.00012†	0.3612
2	Pocket	fabric	fabric	307.3†	0.00012†	0.3612
3	Main Strap	fabric	fabric	307.3†	0.00012†	0.3612
4	Pocket	fabric	fabric	307.3†	0.00012†	0.3612
5	Other	metal	metal	7850	196.5	0.29
6	Main Strap	fabric	fabric	307.3†	0.00012†	0.3612
7	Main Part	metal	metal	7850	196.5	0.29

37. Handbag: e8fe8902cddf44eda16bd7113efbac7e (8 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Main Handle	rubber	rubber	1000	0.0088	0.49
1	Main Strap	fabric	fabric	307.3†	0.00012†	0.3612
2	Main Part	metal	metal	7850	196.5	0.29
3	Pocket	fabric	fabric	307.3†	0.00012†	0.3612
4	Pocket	fabric	fabric	307.3†	0.00012†	0.3612
5	Lock	metal	metal	7850	196.5	0.29
6	Pocket	fabric	fabric	307.3†	0.00012†	0.3612
7	Main Strap	fabric	fabric	307.3†	0.00012†	0.3612

38. Headphone: 9f469c795db945d7a9072d1287dcdf0d (7 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Ear Pad	fabric	fabric	307.3†	0.00012†	0.3612
1	Ear Cup Frame	metal	metal	7850	203	0.29
2	Ear Cup Connector	metal	plastic	1200	2.25	0.3643
3	Headband	metal	metal	7850	203	0.29
4	Ear Pad	fabric	fabric	307.3†	0.00012†	0.3612
5	Ear Cup Frame	metal	metal	7850	203	0.29
6	Ear Cup Connector	metal	plastic	1200	2.25	0.3643

39. Headphone: b4f4aea1dd1342b89cc7347f38776290 (8 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Headband	metal	leather	860	0.0225	0.43
1	Other	plastic	plastic	1200	2.25	0.3643
2	Other	plastic	plastic	1200	2.25	0.3643
3	Other	plastic	glass	2250	65.7	0.228
4	Ear Pad	fabric	fabric	307.3†	0.00012†	0.3612
5	Ear Cup Frame	plastic	metal	7850	203	0.29
6	Ear Pad	fabric	fabric	307.3†	0.00012†	0.3612
7	Ear Cup Frame	plastic	metal	7850	203	0.29

40. Keyboard: fef698ef0b0c4d25918d0e5a700a85ef (4 parts, gallery view 5)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
10	Support Foot	plastic	plastic	1200	2.25	0.3643
11	Support Foot	plastic	plastic	1200	2.25	0.3643
4	Integrated Key	plastic	plastic	1200	2.25	0.3643
9	Casing	metal	plastic	1200	2.25	0.3643

41. Keyboard: fb2a49321c5b48468a8f0fc1233fe207 (11 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
12	Casing	metal	metal	7850	196.5	0.29
13	Keycap	plastic	plastic	1200	2.25	0.3643
14	Keycap	plastic	plastic	1200	2.25	0.3643
15	Keycap	plastic	plastic	1200	2.25	0.3643
16	Keycap	plastic	plastic	1200	2.25	0.3643
17	Keycap	plastic	plastic	1200	2.25	0.3643
18	Keycap	plastic	plastic	1200	2.25	0.3643
19	Keycap	plastic	plastic	1200	2.25	0.3643
20	Keycap	plastic	plastic	1200	2.25	0.3643
21	Keycap	plastic	plastic	1200	2.25	0.3643
22	Keycap	plastic	plastic	1200	2.25	0.3643

42. Knife: c87b4663e6984cd78a5cc8072c8e13b5 (4 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Blade	metal	metal	7850	203	0.29
1	Guard	metal	metal	8600	110	0.33
2	Pommel	metal	metal	8600	110	0.33
3	Grip	metal	rubber	1000	0.0088	0.49

43. Knife: 4a689500b856414da8dd49d7be43f30c (7 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Blade	metal	metal	7850	203	0.29
1	Guard	metal	metal	7800	196.5	0.29
2	Guard	metal	metal	7850	196.5	0.29
3	Guard	metal	metal	7800	196.5	0.29
4	Pommel	metal	metal	7800	196.5	0.29
5	Grip	rubber	rubber	1000	0.0088	0.49
6	Grip	metal	rubber	1000	0.0088	0.49

44. Lamp: 50d4f22f842d400d8cb98b039c354cdc (8 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Shade Cover	fabric	glass	2230	62.6	0.228
1	Bulb	glass	glass	2230	62.6	0.228
2	Socket	plastic	plastic	1200	2.25	0.3643
3	Shade Connection Rod	metal	metal	7850	203	0.29
4	Wire	metal	metal	7850	203	0.29
5	Plug	plastic	metal	7850	203	0.29
6	Switch	plastic	plastic	1200	2.25	0.3643
7	Other	metal	metal	7850	196.5	0.29

45. Lamp: b062b64013604e478d1e4182016caaac (4 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Base	ceramic	metal	8600	110	0.33
1	Switch	plastic	metal	8600	110	0.33
2	Other	metal	metal	8600	110	0.33
3	Shade Cover	fabric	fabric	$307.3^\dagger$	$0.00012^\dagger$	0.3612

46. Laptop Computer: 9b45e7a38db844bd939dbc0d5ea19737 (7 parts, gallery view 1)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Screen	glass	glass	2250	65.7	0.228
1	Camera	plastic	plastic	1200	2.25	0.3643
2	Casing	metal	metal	2710	71	0.33
3	Backplate	wood	metal	7850	196.5	0.29
4	Casing	metal	metal	2710	71	0.33
5	Touchpad	plastic	plastic	1200	2.25	0.3643
6	Buttons	plastic	plastic	1200	2.25	0.3643

47. Laptop Computer: 2e02c7cee4984148bb78269b285425b2 (4 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Casing	metal	metal	2710	71	0.33
1	Keys	plastic	plastic	1200	2.25	0.3643
3	Casing	metal	metal	2710	71	0.33
4	Screen	glass	glass	2250	65.7	0.228

48. Microwave Oven: 1e524fe8e0e546d69baff4d6e20e7533 (9 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Display Screen	plastic	plastic	1200	2.25	0.3643
1	Casing	plastic	metal	7800	196.5	0.29
2	Door Frame	metal	metal	7850	203	0.29
3	Door Handle	metal	metal	7800	196.5	0.29
5	Tray	plastic	plastic	1200	2.25	0.3643
6	Tray	metal	plastic	1200	2.25	0.3643
7	Tray	metal	plastic	1200	2.25	0.3643
8	Tray	metal	plastic	1200	2.25	0.3643
9	Tray	metal	plastic	1200	2.25	0.3643

49. Microwave Oven: be59c580bda64e63bd9e8520f431be55 (6 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Display Screen	plastic	plastic	1200	2.25	0.3643
1	Control Panel Backplate	wood	metal	7800	196.5	0.29
2	Door Glass	glass	glass	2230	62.6	0.228
3	Casing	metal	metal	7800	196.5	0.29
4	Door Handle	wood	metal	7800	196.5	0.29
5	Door Frame	metal	metal	7850	203	0.29

50. Monitor: d8d886b4646145828d62fc49cbd4f8e3 (9 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
1	Base	plastic	plastic	1200	2.25	0.3643
2	Neck	plastic	plastic	1200	2.25	0.3643
3	Ports	plastic	plastic	1200	2.25	0.3643
4	Buttons	plastic	plastic	1200	2.25	0.3643
5	Buttons	plastic	plastic	1200	2.25	0.3643
6	Buttons	plastic	plastic	1200	2.25	0.3643
7	Buttons	plastic	plastic	1200	2.25	0.3643
8	Casing	plastic	plastic	1200	2.25	0.3643
9	Adjustment Mechanism	plastic	plastic	1200	2.25	0.3643

51. Monitor: 000234d39e684deb80b7b8d028681dc1 (4 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Screen	glass	glass	2250	65.7	0.228
1	Casing	plastic	plastic	1200	2.25	0.3643
2	Neck	metal	plastic	1200	2.25	0.3643
3	Base	metal	plastic	1200	2.25	0.3643

52. Mug: 14138954039548a5a6aa49219da9f0c0 (4 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Handle	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
1	Main Body	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
2	Heat Ring	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
3	Heat Ring	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$

53. Padlock: 40bae81948a8425a80aeeaa3904410e6 (9 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Dial	metal	metal	7800	196.5	0.29
1	Dial	metal	metal	8600	127.6	0.27
2	Dial	metal	metal	8600	127.6	0.27
3	Dial	metal	metal	8600	127.6	0.27
4	Reset Hole	metal	plastic	1200	2.25	0.3643
5	Unlock Switch	plastic	plastic	1200	2.25	0.3643
6	Casing	plastic	metal	8600	127.6	0.27
7	U-Shaped Bar	metal	metal	7850	203	0.29
8	End Cap	metal	metal	8600	110	0.33

54. Padlock: 6b97d631e6084f94993bc94a471ca37c (10 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Reset Hole	metal	plastic	1200	2.25	0.3643
1	Unlock Switch	plastic	plastic	1200	2.25	0.3643
2	Dial	metal	metal	7800	196.5	0.29
3	Dial	metal	metal	7800	196.5	0.29
4	Dial	metal	metal	8600	127.6	0.27
5	Dial	metal	metal	8600	127.6	0.27
6	Casing	plastic	metal	8600	127.6	0.27
7	Unlock Switch	plastic	metal	8600	110	0.33
8	U-Shaped Bar	metal	metal	7850	203	0.29
9	End Cap	metal	metal	8600	110	0.33

55. Pen: d1e42a4c243a46598f25779054b2a93d (6 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Clip	metal	metal	7800	196.5	0.29
1	End	metal	plastic	1200	2.25	0.3643
2	Nib	metal	metal	7800	196.5	0.29
3	Tip	metal	metal	7850	203	0.29
4	Grip Area	rubber	rubber	1015	0.000646	0.49
5	Barrel Casing	metal	metal	7850	203	0.29

56. Pen: 5266d47cd31e459887f4e9c68c66e5ad (6 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Nib	metal	metal	7850	196.5	0.29
1	Tip	metal	plastic	1200	2.25	0.3643
2	End	plastic	plastic	1200	2.25	0.3643
3	Cap Body	metal	metal	7850	196.5	0.29
4	Grip Area	rubber	rubber	1015	0.000646	0.49
5	Barrel Casing	plastic	plastic	1200	2.25	0.3643

57. Piano: 78b01b877ba24d9eae4fa2397bb7705c (9 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Pedal	metal	metal	8600	110	0.33
1	Pedal	metal	metal	8600	110	0.33
2	Pedal	metal	metal	8600	110	0.33
3	Pedal Rod	metal	metal	7850	203	0.29
4	Foot	wood	metal	7850	203	0.29
5	Foot	wood	metal	7850	203	0.29
6	Foot	wood	metal	7850	203	0.29
7	Casing	wood	wood	560	9.215	0.3665
8	Keys	plastic	wood	560	9.215	0.3665

58. Piano: d8adb7ed5a744b668965603a73cc0850 (10 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Other	wood	plastic	1200	2.25	0.3643
1	Other	wood	plastic	1200	2.25	0.3643
2	Other	wood	plastic	1200	2.25	0.3643
3	Other	wood	plastic	1200	2.25	0.3643
4	Other	wood	plastic	1200	2.25	0.3643
5	Other	wood	plastic	1200	2.25	0.3643
6	Other	wood	plastic	1200	2.25	0.3643
7	Other	wood	plastic	1200	2.25	0.3643
8	Slider	plastic	plastic	1200	2.25	0.3643
9	Casing	wood	wood	560	9.215	0.3665

59. Pickaxe: b0f929a48c214822a82f541962928d74 (4 parts, gallery view 5)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
3	Pointed End	metal	metal	7850	203	0.29
4	Shaft	wood	wood	560	9.215	0.3665
5	Flat End	metal	metal	7850	203	0.29
6	Connector	metal	metal	7850	203	0.29

60. Pickaxe: 1e2d3b1131b34451846ba35fa76e9d8c (4 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Pointed End	metal	metal	7850	203	0.29
1	Pointed End	metal	metal	7850	203	0.29
2	Connector	metal	metal	7850	203	0.29
3	Grip	rubber	rubber	1000	0.0088	0.49

61. Pliers: 3ec35870870c47efa70978fc9ef9e830 (7 parts, gallery view 1)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Grip Cover	rubber	rubber	1000	0.0088	0.49
1	Central Joint	metal	metal	7850	196.5	0.29
2	Clamping Surface	metal	metal	7850	203	0.29
3	Remaining Jaw Part	metal	metal	7850	203	0.29
4	Grip Cover	rubber	rubber	1000	0.0088	0.49
5	Clamping Surface	metal	metal	7850	203	0.29
6	Remaining Jaw Part	metal	metal	7850	203	0.29

62. Pliers: 36c92a67d5784f1f8e44b2c8c16a1cb3 (7 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Central Pivot	metal	metal	7850	196.5	0.29
1	Grip Cover	rubber	rubber	1000	0.0088	0.49
2	Clamping Surface	metal	metal	7850	203	0.29
3	Remaining Jaw Part	metal	metal	7850	203	0.29
4	Grip Cover	rubber	rubber	1000	0.0088	0.49
5	Clamping Surface	metal	metal	7850	203	0.29
6	Remaining Jaw Part	metal	metal	7850	203	0.29

63. Scissors: 25dbd93da0a74391934251c4a2524c77 (5 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Blade	metal	metal	7850	203	0.29
1	Blade	metal	metal	7850	203	0.29
2	Connector	metal	metal	7850	203	0.29
3	Grip Cover	rubber	rubber	1000	0.0088	0.49
4	Grip Cover	rubber	rubber	1000	0.0088	0.49

64. Scissors: 6c080823d44845a6a3cffec8444af552 (7 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Blade	metal	metal	7850	203	0.29
1	Blade	metal	metal	7850	203	0.29
2	Connector	metal	metal	7850	203	0.29
3	Grip	rubber	metal	7850	203	0.29
4	Grip	rubber	metal	7850	203	0.29
5	Grip Cover	rubber	rubber	1000	0.0088	0.49
6	Grip Cover	rubber	rubber	1000	0.0088	0.49

65. Screwdriver: b5d3b06074a0402e9b13d32c2156d74c (4 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Shaft	metal	metal	7850	203	0.29
1	Handle End	rubber	rubber	1000	0.0088	0.49
2	Connector	metal	metal	7850	203	0.29
3	Grip	rubber	rubber	1000	0.0088	0.49

66. Screwdriver: cfb2020e71fc4662866e0a63b6c9a45c (4 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Shaft	metal	metal	7850	203	0.29
1	Connector	metal	metal	7850	203	0.29
2	Handle End	rubber	rubber	1000	0.0088	0.49
3	Grip	rubber	rubber	1000	0.0088	0.49

67. Shovel: f0fa467580e742da8f783bb5951bcf48 (9 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Surface	metal	wood	560	9.215	0.3665
1	Footrest Area	wood	wood	560	9.215	0.3665
2	Footrest Area	wood	wood	560	9.215	0.3665
3	Connector Body	metal	metal	7850	203	0.29
4	Rod	metal	metal	7850	203	0.29
5	Horizontal Bar	metal	metal	7850	203	0.29
6	Connector Body	metal	metal	7850	196.5	0.29
7	Diagonal Bar	metal	metal	7850	203	0.29
8	Diagonal Bar	metal	metal	7850	203	0.29

68. Shovel: dc2e9e886b104dd1b8754764091e53cb (7 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Horizontal Bar	metal	metal	8600	110	0.33
1	Surface	metal	metal	7850	203	0.29
2	Footrest Area	metal	metal	7850	203	0.29
3	Footrest Area	metal	metal	7850	203	0.29
4	Connector Body	metal	metal	7850	203	0.29
5	Rod	metal	metal	8600	110	0.33
6	Fastening Rivet	metal	metal	7850	203	0.29

69. Skateboard: 2ff01d32fbde4a9da64772864172e3e3 (9 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Rest of Deck	wood	wood	500	7.58	0.3665
1	Baseplate	metal	metal	2710	71	0.33
2	Baseplate	metal	metal	2710	71	0.33
3	Wheel Body	rubber	plastic	1200	2.25	0.3643
4	Wheel Body	rubber	plastic	1200	2.25	0.3643
5	Wheel Body	rubber	plastic	1200	2.25	0.3643
6	Wheel Body	rubber	plastic	1200	2.25	0.3643
7	Kingpin	metal	metal	7850	203	0.29
8	Kingpin	metal	metal	7850	203	0.29

70. Skateboard: 801dfed0f3b04d619fa8cfdad756aeff (12 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Grip Tape	composite	composite	1200	35.41	0.255
1	Rest of Deck	wood	wood	500	7.58	0.3665
10	Wheel Body	rubber	rubber	1180	0.04	0.48
11	Wheel Body	rubber	plastic	1200	2.25	0.3643
2	Baseplate	metal	metal	2710	71	0.33
3	Baseplate	metal	metal	2710	71	0.33
4	Kingpin	metal	metal	7850	203	0.29
5	Kingpin	metal	metal	7850	203	0.29
6	Kingpin	metal	metal	7850	203	0.29
7	Kingpin	metal	metal	7850	203	0.29
8	Wheel Body	rubber	plastic	1200	2.25	0.3643
9	Wheel Body	rubber	plastic	1200	2.25	0.3643

71. Sofa: 58d54c95c3064d51ab1d90edad0644a4 (9 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Seat Cushion	fabric	fabric	307.3†	0.00012†	0.3612
1	Armrest Cushion	fabric	fabric	307.3†	0.00012†	0.3612
2	Armrest Cushion	fabric	fabric	307.3†	0.00012†	0.3612
3	Backrest Cushion	fabric	fabric	307.3†	0.00012†	0.3612
4	Seat Cushion Support Box	wood	wood	560	9.215	0.3665
5	Leg	metal	metal	8600	110	0.33
6	Leg	metal	wood	560	9.215	0.3665
7	Leg	metal	wood	560	9.215	0.3665
8	Leg	metal	wood	560	9.215	0.3665

72. Sofa: 64f9de49a71443e59e0c30d9dc4c1189 (12 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Backrest Cushion	fabric	fabric	307.3†	0.00012†	0.3612
1	Backrest Frame	wood	fabric	307.3†	0.00012†	0.3612
10	Other	wood	fabric	307.3†	0.00012†	0.3612
11	Armrest Support Box	wood	wood	560	9.215	0.3665
2	Other	fabric	fabric	307.3†	0.00012†	0.3612
3	Armrest Support Box	wood	wood	560	9.215	0.3665
4	Seat Cushion	fabric	fabric	307.3†	0.00012†	0.3612
5	Seat Cushion Support Box	wood	wood	560	9.215	0.3665
6	Leg	metal	wood	560	9.215	0.3665
7	Leg	metal	wood	560	9.215	0.3665
8	Leg	metal	wood	560	9.215	0.3665
9	Leg	metal	wood	560	9.215	0.3665

73. Sword: 180656a44ef047e99492550c9d53d221 (4 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Blade	metal	metal	7850	203	0.29
1	Pommel	metal	metal	8600	110	0.33
2	Grip	metal	metal	8600	110	0.33
3	Guard	metal	metal	8600	110	0.33

74. Sword: 259f0c975fdc41388fe1fa6bc0344b4a (7 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Blade	metal	metal	7850	203	0.29
1	Other	metal	metal	7800	196.5	0.29
2	Other	plastic	metal	7800	196.5	0.29
3	Guard	metal	metal	7800	196.5	0.29
4	Guard	metal	metal	7800	196.5	0.29
5	Grip	metal	rubber	1000	0.0088	0.49
6	Pommel	metal	metal	7800	196.5	0.29

75. Table: 1338705a26de4b89808e687e243b46b2 (7 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Surface Panel	wood	wood	560	9.215	0.3665
1	Table Frame	wood	wood	560	9.215	0.3665
2	Leg	wood	wood	560	9.215	0.3665
3	Leg	wood	wood	560	9.215	0.3665
4	Leg	wood	wood	560	9.215	0.3665
5	Leg	wood	wood	560	9.215	0.3665
6	Other	metal	wood	560	9.215	0.3665

76. Table: 78fd469edd0c40d3827241b4e5468615 (4 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Surface Panel	wood	wood	560	9.215	0.3665
1	Surface Panel	wood	wood	560	9.215	0.3665
2	Base	wood	metal	7850	203	0.29
3	Central Support Column	metal	wood	560	9.215	0.3665

77. Teapot: 3f87c660931b4b658f1d83fd00f68ca0 (5 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Handle	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
1	Spout	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
2	Knob	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
3	Body	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
4	Cover	metal	metal	7850	196.5	0.29

78. Teapot: 1082ac22b8634e5ead2f2ec31597420b (5 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Spout	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
1	Body	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
2	Handle	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
3	Knob	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
4	Cover	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$

79. Toaster: 68a31a8eb07d4b84aeb3b13484d4c9e8 (11 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Lever Body	metal	metal	7850	196.5	0.29
1	Slide Rail	metal	plastic	1200	2.25	0.3643
10	Inner Shell	metal	metal	7850	196.5	0.29
2	Horizontal Bar	metal	metal	7850	203	0.29
3	Horizontal Bar	metal	metal	7850	203	0.29
4	Casing	metal	metal	7850	196.5	0.29
5	Bread Slot Side Wall	metal	metal	8600	127.6	0.27
6	Support Foot	metal	rubber	1015	0.000646	0.49
7	Bread Slot Side Wall	metal	metal	8600	127.6	0.27
8	Slide Rail	metal	plastic	1200	2.25	0.3643
9	Lever Body	metal	plastic	1200	2.25	0.3643

80. Toilet: c4807f7c0f2247e697dc93daa915a188 (11 parts, gallery view 6)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Lid	plastic	plastic	1200	2.25	0.3643
1	Seat Body	plastic	plastic	1200	2.25	0.3643
10	Other	plastic	plastic	1200	2.25	0.3643
2	Seat Cushion	plastic	plastic	1200	2.25	0.3643
3	Seat Cushion	plastic	plastic	1200	2.25	0.3643
4	Connection Part	plastic	plastic	1200	2.25	0.3643
5	Connection Part	plastic	plastic	1200	2.25	0.3643
6	Base	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
7	Screw	metal	metal	7850	203	0.29
8	Screw	metal	metal	7850	203	0.29
9	Button	plastic	plastic	1200	2.25	0.3643

81. Toilet: 4d24c6ab618c4754802ffd94de41c0f2 (6 parts, gallery view 5)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Button	plastic	plastic	1200	2.25	0.3643
1	Tank Lid	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
2	Tank Wall	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
3	Lid	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
4	Seat Body	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$
5	Base	ceramic	ceramic	$2350^{\ddagger}$	$75.03^{\ddagger}$	$0.204^{\ddagger}$

82. Umbrella: d9a662106733465f98b6875e8b594aab (5 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Fabric	fabric	fabric	$307.3^{\dagger}$	$0.00012^{\dagger}$	0.3612
1	Tip	plastic	plastic	1200	2.25	0.3643
2	Central Rod	wood	wood	560	9.215	0.3665
3	Central Support Column	metal	wood	560	9.215	0.3665
4	Base	wood	metal	2710	71	0.33

83. Umbrella: ee3be49e9b8e45c8a761616c27cf3797 (7 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Tip	metal	plastic	1200	2.25	0.3643
1	Grip	rubber	rubber	1000	0.0088	0.49
2	Fabric	fabric	fabric	$307.3^{\dagger}$	$0.00012^{\dagger}$	0.3612
3	Stretcher	metal	metal	2710	71	0.33
4	Central Shaft	plastic	plastic	1200	2.25	0.3643
5	Central Rod	metal	metal	2710	71	0.33
6	Rib Rod	metal	metal	2710	71	0.33

84. Washing Machine: 7e088f1294f74a6a81e86309884e1df3 (10 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Frame	metal	metal	7850	203	0.29
1	Door Frame	metal	metal	7850	203	0.29
10	foot	metal	metal	7850	203	0.29
2	Inner Drum	metal	metal	7850	196.5	0.29
3	Drum Shaft	metal	metal	7850	196.5	0.29
5	Frame	metal	metal	7850	203	0.29
6	Knob	metal	metal	7850	196.5	0.29
7	foot	metal	metal	7850	203	0.29
8	foot	metal	metal	7850	203	0.29
9	foot	metal	metal	7850	203	0.29

85. Watch: 0280fb9e62e142229e82ea2c77ffb1e4 (6 parts, gallery view 3)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Crystal	glass	glass	2250	65.7	0.228
1	Stem	metal	plastic	1200	2.25	0.3643
2	Crown Wheel	metal	metal	7850	196.5	0.29
3	Case Back	metal	metal	7850	196.5	0.29
4	Bezel	metal	metal	7850	203	0.29
5	Middle Case	metal	metal	7850	196.5	0.29

86. Watch: d6ea29a09e5e45fe9b2c335460559d73 (11 parts, gallery view 2)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Strap Body	metal	leather	860	0.0225	0.43
1	Strap Body	metal	leather	860	0.0225	0.43
10	Center Axle	metal	metal	7800	196.5	0.29
2	Keeper	metal	metal	7850	196.5	0.29
3	Crown Wheel	metal	metal	7800	196.5	0.29
4	Middle Case	metal	metal	7800	196.5	0.29
5	Case Back	metal	metal	7800	196.5	0.29
6	Surface	metal	metal	7800	196.5	0.29
7	Hands	metal	metal	7800	196.5	0.29
8	Hands	metal	metal	7800	196.5	0.29
9	Hands	metal	metal	7800	196.5	0.29

87. Wrench: 158a5e095e2d4161b8f440426ea24eea (7 parts, gallery view 1)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Grip	rubber	rubber	1000	0.0088	0.49
1	Adjustment Screw	metal	metal	7850	203	0.29
2	Jaw	metal	metal	7850	203	0.29
3	Connecting Rod	metal	metal	7850	203	0.29
4	Jaw	metal	metal	7850	203	0.29
5	Connector	metal	metal	7850	196.5	0.29
6	Handle Rod	metal	metal	7850	203	0.29

88. Wrench: 874f981eab3e42d49e30411a7fc5f10a (4 parts, gallery view 4)

ID	Part	Released	Predicted class	ρ (kg/m^3)	E (GPa)	ν
0	Adjustment Screw	metal	metal	7850	203	0.29
1	Adjustable Jaw	metal	metal	7850	203	0.29
2	Fixed Jaw	metal	metal	7850	203	0.29
3	Grip	metal	metal	7850	196.5	0.29