

Fantastic Breaks: A Dataset of Paired 3D Scans of Real-World Broken Objects and Their Complete Counterparts

Nikolas Lamb, Cameron Palmer, Benjamin Molloy, Sean Banerjee, Natasha Kholgade Banerjee
Clarkson University, Potsdam NY, USA

{lambne, campalme, molloybr, sbanerje, nbanerje}@clarkson.edu

Abstract

Automated shape repair approaches currently lack access to datasets that describe real-world damaged geometry. We present Fantastic Breaks (and Where to Find Them: <https://terascale-all-sensing-research-studio.github.io/FantasticBreaks>), a dataset containing scanned, waterproofed, and cleaned 3D meshes for 150 broken objects, paired and geometrically aligned with complete counterparts. Fantastic Breaks contains class and material labels, proxy repair parts that join to broken meshes to generate complete meshes, and manually annotated fracture boundaries. Through a detailed analysis of fracture geometry, we reveal differences between Fantastic Breaks and synthetic fracture datasets generated using geometric and physics-based methods. We show experimental shape repair evaluation with Fantastic Breaks using multiple learning-based approaches pre-trained with synthetic datasets and re-trained with subset of Fantastic Breaks.

1. Introduction

Damage to objects is an expected occurrence of everyday real-world usage. However, when damage occurs, objects that could be repaired are often thrown out. Additive manufacturing techniques are rapidly becoming accessible at the consumer level, with 3D printing technologies available for materials such as plastics, metals, and even ceramics and wood. Though current approaches for repair have been largely manual and restricted to niche areas such as cultural heritage restoration, a large body of recent research has emerged on the automated reversal of damage, including reassembly of fractured parts using 3D scans [4, 9, 17, 19, 21–23, 38, 39, 52, 54, 55], or generation of new repair parts when portions of the original object are irretrievably lost due to damage [18, 20, 28–30, 32, 35, 40, 47, 48]. Geometry-driven approaches based on shape matching [4, 17, 18, 21–23, 27, 32, 35, 38–40, 47, 48] that are not usable for objects of unknown complete geometry have given




Figure 1. We present *Fantastic Breaks*, a dataset of 3D scans of real-world broken objects (top) aligned to 3D scans of complete counterparts (bottom). Objects span classes such as mugs, plates, statues, jars, and bowls—household items prone to damage.

way to learning-driven approaches [9, 19, 20, 28–30, 52, 54] aimed at generalization to repair at a large scale.

However, a principal challenge limiting understanding of real-world damage and limiting application-focused evaluation of repair approaches is that datasets of real-world damage for consumer space objects are virtually non-existent. Current learning-driven approaches for repair use datasets where fracture-based damage is synthetically generated using geometric approaches such as Boolean operations with primitives [9, 18, 28–31]. As they make assumptions about the fracture process rather than using data-driven fracture generation, the practical usability of such geometric methods for repairing real fractures is unknown. With *Breaking Bad*, Sellán et al. [44] have taken a first step toward large-scale fracture dataset generation. *Breaking Bad* consists of 3D shapes from Thingi10k [56] and PartNet [37] subjected to physics-based damage using fracture modes [45]. By removing macro-scale shape assumptions embodied by geometric primitives, *Breaking Bad* is a promising step for research in shape assembly and repair. However, fractures in *Breaking Bad* suffer from typical issues of resolution and simulation time step size that underlie physics simulations.

The dataset is thus unfortunately hindered in providing a faithful representation of real-world damage.

In this work, we contribute *Fantastic Breaks*, the first dataset of 3D scans of damaged objects paired with 3D scans of their complete non-damaged counterparts, as shown in Figure 1. Each damaged object—hereafter referred to a broken object due to the nature of damage suffered—is 3D scanned and geometrically registered to a scan of the undamaged complete object, such that intact regions of the complete and broken scans are aligned. At the time of publication, the dataset contains 150 broken/complete pairs.

Our damage infliction often leaves one broken part intact, destroying or over-fragmenting the remainder of the object. We use an off-the-shelf subtraction-based approach [27] to generate repair part proxies from the intact broken part aligned to its complete counterpart. We also provide manually annotated class labels, material labels, and annotated fractured regions on the broken meshes. Our work emulates endeavors of groups in vision and robotics that contribute datasets of 3D scanned everyday-use objects [5, 6, 13, 25, 46]. Our analysis of real-world fracture properties in *Fantastic Breaks* reveals fine-scale fracture structure that enables our dataset to overcome the drawbacks of synthetic datasets. We use our dataset to evaluate existing shape repair approaches on real fractured objects.

We summarize our contributions as follows:

1. We contribute the first 3D scanned real-world dataset of geometrically aligned broken/complete object pairs to enable application-focused evaluation of repair.
2. We provide class, material, and fracture surface annotations, and ground truth repair part proxies.
3. We contribute a geometric analysis of *Fantastic Breaks* in comparison to existing synthetic fracture datasets.
4. We provide evaluations of existing shape repair approaches using *Fantastic Breaks*.

2. Related Work

Fracture Datasets. Existing real-world fracture collections are restricted to scanning of multiple shards corresponding to a small set of originally intact objects, e.g., 7 objects [23], 3 frescoes [4], and 3 large-scale structures (Akrotiri settlement, Tongeren Roman excavation, and one fresco) [17]. Since fracture acquisition is performed post-damage for historical objects where no known counterparts exist, the datasets lack knowledge of the complete proxies. The Hampson Museum cultural heritage dataset [42] contains 3D scans for 138 cultural heritage objects. The dataset lacks paired damaged/complete data, or annotations to reveal what objects are damaged and what are intact, preventing them from being used to train repair or assembly approaches. Hong et al. [22] show assembly results using a dataset of 5 shattered pots, with 3D scans acquired for the shattered fragments and the original pots. The origi-

nal scans are used as evaluation oracles. Lamb et al. [27] repair 22 damaged objects by subtracting the damaged objects from *a priori* known complete proxies.

Recognizing the need for large-scale datasets for learning-driven repair, a few datasets contain synthetic or scanned models subjected to synthetic fracture using geometric techniques such as subtracting primitives [9, 18, 28–31], or using physics models of fracture [44]. As we demonstrate in Section 4, real-world physical damage demonstrates geometric characteristics that differ from the break patterns of synthetically generated damage. Geometric fracture models [9, 18, 28–31] are only as precise as the primitive being used for fracture, e.g., Chen et al. [9] use five simplistic cut functions—planar, sine, parabolic, square, and pulse—modeled as oriented height fields. Lamb et al. [28–31] subtract randomly rotated and translated geometric primitives such as a cube, icosphere, and sub-divided icosphere, with and without random surface perturbation to simulate micro-scale detail. Gregor et al. [18] use spheres whose surfaces are perturbed using small-scale details of a single digitized material. In all cases, mid-scale detail is represented as cutouts via analytical primitives, which is not generalizable to real-world damage.

Though approaches exist to perform fracturing of single objects using physics simulations [10, 14, 34, 50, 51], which could theoretically be used to generate a fractured object dataset, they suffer from drawbacks that prevent them from being used at scale. (1) *Unrepairable Objects*. Objects are repairable only if a large part remains intact. Approaches that cannot control the number of parts [10, 14, 50, 51] may require many trials to obtain a repairable object. (2) *Numerical Instability*. Polygon division may cause simulation failure due to roundoff error [10, 34]. (3) *Limited Diversity*. Some approaches tend to produce planar fractures [14] or limit crack speed to ensure PDE convergence [50]. (4) *Runtimes*. Fractures on objects with similar degrees of freedom to ours (2.5M) require multiple hours, e.g. 6h/1M [14], 7h/2.5M [50], 3.5h/2.5M [34], 26h/2.5M [10], 4h/1M [51]. (5) *Parameter Tuning*. Physics parameters, e.g. density, fracture threshold, and timestep size [50], need tuning to produce diverse fractures, a process which has a “notable learning curve” [51]. In summary, excessive parameter tuning and long runtimes make dataset creation using physics simulations intractable. The only physics-based fractured dataset is Breaking Bad [44]. The Breaking Bad dataset of Sellán et al. uses their earlier work [45] to model and simulate fracture modes of the object. Their approach struggles to generate high-resolution detail, as it requires evolving the simulation over small time steps that can be computationally infeasible at a large scale. Our *Fantastic Breaks* dataset fills the gap in the lack of real-world damaged object datasets by contributing an even-now growing repository of 3D scanned real-world damaged objects rigidly aligned

with 3D scans of their real-world counterparts.

Real-World Object Datasets. Capture of large-scale complete 3D scans of objects is an arduous task, due to the need for a multi-staged approach consisting of multiple presentations of the object to a scanner, registration of scans, cleaning to correct imprecise geometry, eliminate holes, and correct deep concavities, and, depending on the application, waterproofing to ensure closed surfaces. Massive 3D datasets typically comprise single-viewpoint RGB-D images [24, 26, 53], room-scale 3D scans [12], or scans of large objects such as motorcycles, large statues, and benches [11] that cannot be readily impacted to capture real-world damage. The Berkeley BigBIRD dataset [46] contained 100 3D models at publication time and has grown to 125 models. Objects were recorded on a turntable and models were created by stitching multiple Kinect and Canon camera images. Using the BigBIRD setup, Calli et al. [5, 6] collected the YCB dataset for a set of 77 objects as a collaboration between Yale, CMU, and UC Berkeley. Due to placement on a turntable, hidden portions, e.g., the base or concavities not visible to the cameras, are not captured. The Karlsruhe Institute of Technology (KIT) dataset [25] consists of 145 objects captured using a Konica Minolta Vi-900 digitizer, with multiple scans acquired to capture object bases. With 300 3D scans (150 broken, 150 paired complete counterparts, and growing), *Fantastic Breaks* is comparable to these prior tabletop datasets [5, 6, 25, 46] in terms of total items.

Large-scale table-top 3D scanned object datasets are few in number. Google Scanned Objects (GSO) [13] consists of 1,030 3D scans of household objects. GSO is collected by 3D imaging projector-cast patterns through a collaboration between Google and Open Robotics. Another large-scale dataset, AKB-48 [33], contains 2,037 3D scans of complete objects. Scans are obtained using an Einscan 3D scanner and an Intel RealSense depth camera. Unlike *Fantastic Breaks*, neither GSO nor AKB-48 contain broken objects.

Within the larger problem domain of minimizing waste, a few 2D datasets have arisen for object detection and segmentation in waste images [2, 43, 49]. Object materials span cardboard, plastic, glass, and metal. An opportunity exists to capture RGB-D images of waste, identify objects capable of being repaired, and geometrically couple them with our dataset to conduct in-the-wild repair.

Shape Repair. The creation of the *Fantastic Breaks* dataset is motivated by applications in object repair. When object parts are available, shape assembly approaches focus on joining 3D shape representations of the object parts. Geometric approaches exist to match fracture boundaries via segmentation [38, 39], feature description extraction and geometric model refinement [23], align fractured shapes to a proxy template [55], or conduct iterative registration similar

to structure from motion [21, 22]. Several learning-based approaches exist to provide assembly, most of which assume holistic parts with simple surfaces [19, 52, 54], though one assumes arbitrary geometry at shape boundaries [9].

When parts of an object are irretrievable, shape repair approaches address generation of the lost parts. Early automation approaches to circumvent historically manual repair included finding symmetries and self-similarities in objects [18, 32, 35, 40, 47, 48], though these approaches are unsuccessful when non-symmetric object parts are broken off. One approach [27] circumvents the small-scale artifact issue of Boolean subtraction by automatically extracting and joining exterior and fractured regions for repair parts using real-world damaged and complete scans. The approach requires the scan of a complete 3D proxy to be provided as input, that may not be feasible for one-of-a-kind instances, or even if available, may prove tedious to obtain.

Recent work such as 3D-ORGAN [20], MendNet [30], DeepMend [29], and DeepJoin [28] uses deep learning to conduct shape repair without knowledge of the complete proxy by representing damaged, complete, and repair shapes using voxels [20] or deep functions [28–30]. DeepMend and DeepJoin report higher success due to the use of implicit functions which enable representation of arbitrary resolutions, and due to the expression of fractured and restoration shapes in terms of constructive solid geometry operations between the complete object and a break surface whose representation is learnt during training. Though DeepMend and DeepJoin show qualitative results on a few examples of 3D scans of real-world damaged objects, in all cases, training and quantitative evaluation is conducted using datasets with synthetically-generated fractures.

3. Data Collection and Processing

Object Acquisition. We conducted a community-wide acquisition of everyday household objects that suffer damage. Our goal was to have a collection consisting not only of objects that can be damaged, but also of objects that have already suffered damage. To perform our acquisition, we made purchases at the local thrift store where we found intact and damaged objects, and requested donations from the local community. Though the *Fantastic Breaks* dataset as presented in this paper contains damaged/complete pairs, we have also collected pre-damaged objects that may lack complete counterparts, as they provide insight into real-world fracture. Our collection consists of commonly damaged household objects such as mugs, plates, and statues spanning materials such as ceramics, plastics, glass, and wood. In some instances, we were able to pre-acquire pairs, where one object in the pair was damaged and the other was intact. In most cases, we manually damaged a complete object to obtain the broken version. When possible, we attempted to acquire pairs of complete objects, and dam-



Figure 2. Left: Complete mug on top, and broken mug on the bottom showing main intact part and the shattered handle. Right: Object on scanner with 3D scan shown in the inset.



Figure 3. Example broken objects on top and broken/complete pairs on the bottom acquired to build the *Fantastic Breaks* dataset.

age one of the objects, enabling us to store physical damaged/complete pairs. We inflicted damage by dropping the object, striking the object with a rubber mallet or metal hammer, and, in 3 cases, snapping the object after anchoring it against a table. Figure 2 shows an example mug with the handle shattered after striking with a mallet. Figure 3 shows broken objects and broken/complete pairs in our dataset.

3D Scanning. We use an Einscan SP 3D turntable-based scanner to acquire 3D scans. We use Einscan’s proprietary software, EXScan, to operate the scanner. The scanning operation rotates the scanner’s turntable 8 times, acquires 8 2.5D images of the object via an attached RGB-D sensor, and fuses the images into a 3D mesh. Given the diversity of object geometry where objects may contain deep concavities or complex fractures, we had to conduct careful staging of each object and perform multiple presentations of the object to maximize acquisition of the object surface. The presentation count ranged from 2 for flat objects such as plates or objects with higher convexity such as statues, to 6 for objects with concavities such as mugs. We used the registration tool in EXScan to fuse scans into a 3D mesh.

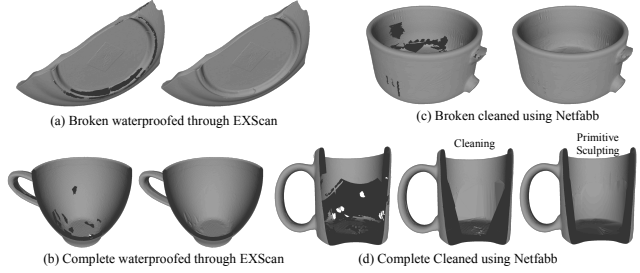


Figure 4. Meshes prior to (left) and after cleaning and waterproofing for (a) broken and (b) complete object examples waterproofed using EXScan, and (c) broken and (d) complete objects cleaned using Netfabb. For (d), we use a primitive to fix the erroneously angled surfaces created by Netfabb mesh repair tools. Complete objects are shown as cutaways to reveal pre-cleaning artifacts.

Mesh Cleaning. Given the 3D scan of a complete or broken object, we employ a sequence of operations to ensure that the meshes are of high quality. We visually inspect each model, and if the model shows any small holes in the mesh, we use the waterproofing tool built into EXScan to ensure that each mesh is a closed 2D surface. While performing waterproofing in EXScan is preferred, if the model contains large holes or artifacts, we repair the mesh by performing manual hole-filling using Autodesk Netfabb. For objects with deep concavities such as cups and mugs, the concavities may not have been well-presented during the scanning process, in which case the interior regions may not be metrically accurate. If possible, we repair the mesh using subtraction with a geometric primitive, e.g., a cylinder. We apply the “Extended Repair” scripts in Netfabb to merge nearby vertices, remove double and flipped triangles, close all holes, wrap the mesh to remove interior faces, and remove small connected components. Figure 4 demonstrates examples of cleaned meshes.

Mesh Orientation. After cleaning, we manually orient each mesh such that its principal axes are aligned to the Cartesian axes. We align the base of the object with the xz -plane, and we rotate the object about the y -axis to ensure alignment within its category. For example, we align mugs such that the handles are aligned with the negative x direction, and statues so that they face in the negative x direction.

Mesh Alignment. Given the cleaned meshes for the broken object and its complete counterpart, we transform the broken mesh such that its intact regions are aligned with the corresponding regions of the complete mesh. We perform an initial manual alignment of the broken mesh to the complete mesh, and refine the alignment using the iterative closest point (ICP) [3] algorithm. We conduct a post-alignment normalization of each mesh to ensure similarly scaled data for learning-driven repair. We perform normalization by scaling the broken and complete mesh such that the complete mesh is entirely contained within a unit cube. We

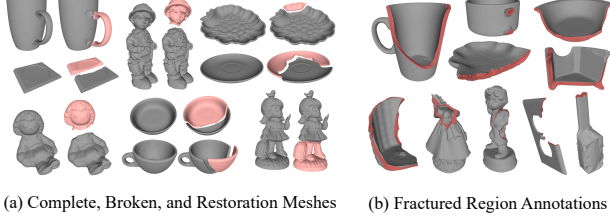


Figure 5. (a) For each triplet of meshes, we show the complete mesh on the left, the broken mesh aligned to the complete mesh on the right, and the restoration in pink detached from the broken. (b) Broken meshes (gray) with labeled fracture surface (red).

provide access to non-normalized and normalized meshes as part of the dataset. Figure 5(a) shows examples of the broken mesh aligned to the complete mesh.

Ground Truth Restoration Estimation. Our data collection involves scanning of broken objects that have either been acquired as pre-damaged, or have had damage inflicted through a destructive fracture process, such that only one part of the object remains intact. Such an occurrence is not uncommon, e.g., as shown in Figure 2, the fracture process causes the handle to further fragment into a number of small pieces that are infeasible to reassemble. However, shape repair approaches can benefit from geometric knowledge of ground truth parts needed to repair the object, in order to perform training and evaluation. We contribute proxy ground truth 3D meshes for the repair parts. We synthetically generate the repair part proxies using the approach of Lamb et al. [27]. Given aligned complete and broken meshes, the approach recovers restoration meshes that smoothly join the broken shape to repair the object, lack small-scale artifacts caused by Boolean subtraction, and lack grooves at the fracture boundary demonstrated by approximate subtraction techniques based on distance thresholding. Figure 5(a) shows restoration meshes generated for cleaned and waterproofed broken and complete scans.

Ground Truth Fracture Surface Annotation. To assist with approaches that rely on accurate knowledge of the fracture surface for training or evaluation, we manually annotate triangles corresponding to the fracture surface. Figure 5(b) shows example ground truth fracture annotations.

4. Analysis of *Fantastic Breaks*

Summary of Dataset. We have annotated each object with a category level, its material, and the approach of damaging the object. At the time of publication we have acquired 214 and 241 physical damaged and complete objects, and 195 and 218 damaged and complete object scans. Among the 195 broken objects, 15 objects are naturally damaged, and the remaining have damage inflicted manually by dropping (26), striking using a mallet (126), strik-

Table 1. Object distribution by class (Phys.=Physical).

	Class	Mug	Plate	Statue	Bowl	Cup	Jar	Coaster	Box	Misc	Total
Phys.	Broken	46	43	39	23	12	13	6	6	26	214
	Complete	42	45	43	27	16	19	7	7	35	241
Scan	Broken	43	41	36	21	10	12	6	5	21	195
	Complete	35	44	41	24	14	15	7	7	31	218
	Pairs	30	35	30	17	8	6	6	3	15	150

Table 2. Object distribution by material (Mat., Phys.=Physical).

	Mat.	Ceramic	Plastic	Glass	Plaster	Wood	Other	Total
Phys.	Broken	165	21	10	4	4	10	214
	Complete	170	30	10	3	8	20	241
Scan	Broken	155	17	6	4	4	9	195
	Complete	151	28	9	3	8	19	218
	Pairs	115	13	7	4	3	8	150

ing using a hammer (25), and anchoring and snapping (3). 150 of the broken objects have a counterpart within the 241 complete objects. We have acquired pre-cleaned scanned meshes for 195 of the broken objects and 218 of the complete objects. 150 of the 214 broken scanned objects are paired with their respective 150 complete scanned objects. These 150 pairs are fully cleaned and have restorations extracted. Tables 1 and 2 summarize physical, scanned, and paired scanned broken and complete objects by class and material respectively. This is a growing dataset, i.e., the physical and scanned sets are continuing to expand. We expect to have all 241 complete objects broken, and we continue to process complete and broken objects.

Analysis of Geometric Properties. We quantitatively compare geometric properties of the *Fantastic Breaks* dataset to the *everyday* subset of *Breaking Bad* [44], and the Geometric Breaks dataset provided by Lamb et al. [28], both of which contain synthetic fractures. The *everyday* subset of *Breaking Bad* contains 542 objects each fractured 100 times, for a total of 54,200 broken objects. Breaks are obtained using a physics driven fracturing technique in which a set of fracture modes are computed and used to simulate object fracture patterns that result from an impact to the object. The Geometric Breaks dataset contains 24,208 objects from ShapeNet [7] and 1,042 objects from the GSO dataset. Objects are fractured by subtracting them with a randomized convex geometric primitive. For *Breaking Bad* we compute quantitative metrics over the 12,939 *everyday* objects that contain 2 broken parts. For Geometric Breaks, we compute quantitative metrics over all 25,250 objects. For *Fantastic Breaks*, we compute quantitative metrics over the 150 paired broken objects.

We summarize the mean number of faces and vertices for broken meshes in each dataset on the left of Table 3. Our broken meshes are, on average, at least ten times more

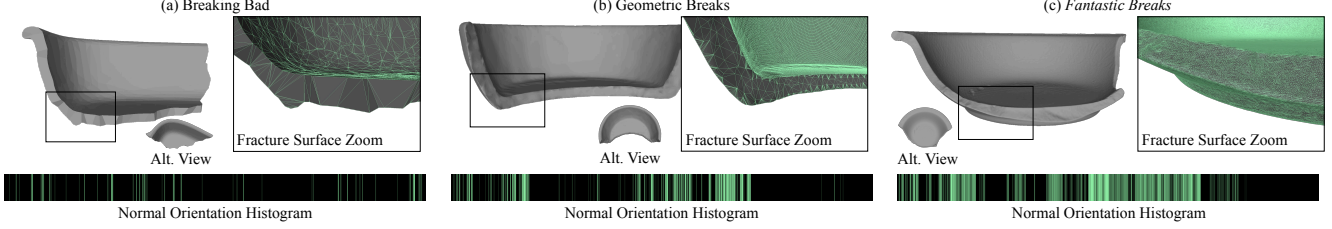


Figure 6. Broken shapes with inset showing geometric detail, alternate (alt.) top-down views, and normal orientation histograms (NOHs) for (a) Breaking Bad, (b) Geometric Breaks, and (c) *Fantastic Breaks*. NOHs are viewed as images where higher intensity lines represent higher bin counts. (a) Breaking Bad shapes show sparse structure due to limitations on fracture resolution. (b) While Geometric Breaks shapes have semi-dense structure, they demonstrate regularity attributed to the breaking primitive. (c) *Fantastic Breaks* shapes have dense surface structure and irregular break patterns characteristic of arbitrary real-world fracture.

Table 3. Number of vertices, faces, and convexity percentiles for Breaking Bad, Geometric Breaks, and *Fantastic Breaks* datasets. Max Vertices and Faces and min Convexity values are bolded.

Dataset	# Vertices	# Faces	C 25th	C 50th	C 75th
Breaking Bad	4,578	17,876	0.286	0.536	0.828
Geometric	52,719	105,436	0.229	0.480	0.746
<i>Fantastic</i>	566,630	1,133,837	0.238	0.320	0.519

dense than existing fractured object datasets. As observed by Sellán et al. [44], an indicator of fractures generated using pre-computed fracture methods are a large number of convex fractured shapes. In the final three columns of Table 3 we provide the 25th, 50th, and 75th percentiles of broken shape convexity. We compute convexity as the ratio of the volume of the broken mesh to the volume of the mesh’s convex hull, as used by Attene et al. [1]. Highly convex shapes show a convexity value near 1. Our dataset shows a consistently lower convexity than other datasets, with a 75th convexity percentile of 0.519, compared to 0.828 for Breaking Bad and 0.746 for Geometric Breaks.

Though synthetically generated fractured objects may show coarse geometric variation, they struggle to generate fine detail at the fractured region. As shown in Figures 6(a) and 6(b), synthetically generated fractured objects are characterized by piecewise fractured regions that lack high frequency surface variation, i.e. they do not capture the fine-scale variability of real fractures. Simulation of object fractures with fine-scale surface variability of the same resolution as real object fractures using a physics engine as done in Breaking Bad is intractable with current hardware. Fractures generated by subtracting geometric primitives demonstrate unnatural regularity as shown in Figure 6(b).

To quantify coarse and fine-scale surface variability, we introduce the spherical occupancy metric, which measures the degree to which normals on the fracture surface occupy the space of all possible orientations when grouped into n discrete bins. For a given broken object, we extract the

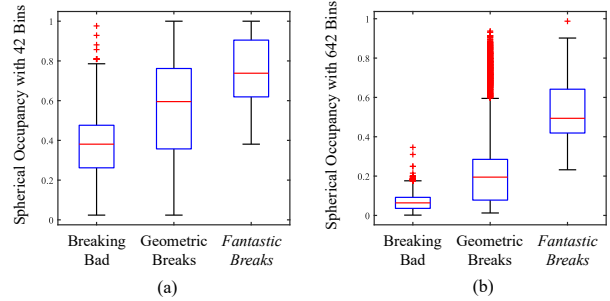


Figure 7. Box plots for distribution of spherical occupancy over objects from Breaking Bad, Geometric Breaks, and *Fantastic Breaks* using (a) 42 bins and (b) 642 bins.

fractured region of the broken mesh. We obtain n evenly spaced points on the unit sphere, and bin the normal vectors of the fractured region into one of n bins based on their closest point on the unit sphere to generate a normal orientation histogram (NOH) for the object. We show NOHs for three objects in Figure 6. To compute spherical occupancy, we count the number of occupied bins in the NOH and divide by n . A fractured region with coarse geometry is expected to show low spherical occupancy, as its normals span a small portion of the space of all possible orientations. The choice for n determines the scale of surface variability measured. A small value for n measures coarse surface variability, i.e. highly concave or convex fractures will have a high spherical occupancy. A large value for n measures fine-scale surface variability. If n is large, only objects with fine-scale surface variability will show a high spherical occupancy, as the space of orientations is larger.

We show quantitative results for the spherical occupancy for two values of n in Figure 7(a) and Figure 7(b). To measure coarse and fine surface variability, we set n to 42 and 642 respectively. As shown in Figure 7(a), most objects in *Fantastic Breaks* have a high spherical occupancy when $n = 42$, indicating that objects in *Fantastic Breaks* have a high degree of coarse surface variability. *Fantastic Breaks*

shows the highest mean spherical occupancy value of 0.747 when n is 42, compared to 0.367 for Breaking Bad and 0.569 for Geometric Breaks. *Fantastic Breaks* also contains mostly objects with high fine-scale surface variability, as shown by the lower spread in Figure 7(b). *Fantastic Breaks* shows the highest median spherical occupancy when n is 642, indicating a large number of fractures with fine-scale surface variability, and includes the object with the highest spherical occupancy across all datasets. In contrast, the piecewise fractures of Breaking Bad and Geometric Breaks, demonstrate consistently lower occupancy, as shown in Figure 7(b), indicating a lack of fine-scale surface variability. *Fantastic Breaks* also has the highest mean spherical occupancy of 0.525 when $n = 642$, compared to 0.066 for Breaking Bad and 0.204 for Geometric Breaks.

5. Experimental Evaluation

As we provide complete, broken, and restoration shapes with ground truth fractured region annotations, our dataset may be used to train and evaluate approaches that perform automated shape repair. We test three prior shape repair approaches on our dataset: MendNet [30], DeepMend [29], and DeepJoin [28]. These approaches generate repair parts assuming that the missing part has been lost or destroyed during the damage process. MendNet, DeepMend, and DeepJoin represent shapes by learning a function to reconstruct shapes as implicit surfaces, and require watertight 3D meshes as input. DeepJoin also requires meshes to have surface normals, which *Fantastic Breaks* provides.

For shape repair methods, we pre-train a given network on the Geometric Breaks dataset or Breaking Bad dataset. After training for 2,000 epochs on synthetically fractured objects, we train for an additional 1,000 epochs on objects from the *Fantastic Breaks* dataset. To generate training data, for each broken, complete, and restoration mesh tuple in each dataset, we sample points on the surface of each mesh and compute the signed distance function (SDF), occupancy, and normal field value for each point. For DeepMend and DeepJoin we compute a break surface that acts as a proxy for the fracture by fitting a thin-plate spline to the fracture vertices, as described by Lamb et al. [28]. We train with 5,068 objects from the mugs, jars, bottles, and GSO classes from Geometric Breaks, and hold out the remaining 45 for evaluation. For Breaking Bad, we train with everyday objects that have 2 parts. Breaking Bad objects are not waterproof, so we perform waterproofing before training. We train with 2,575 objects from the everyday subset of Breaking Bad that are successfully waterproofed, and hold out the remaining 45 for evaluation. *Fantastic Breaks* has 105 train and 45 test objects.

Figure 8 shows repairs for broken objects before and after re-training on *Fantastic Breaks*. Before re-training on real objects, DeepMend and DeepJoin may struggle to gen-

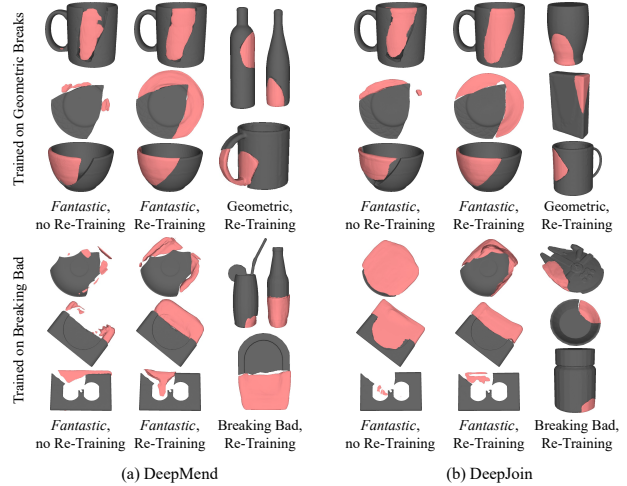


Figure 8. Predicted repair shapes in pink trained with (top) Geometric Breaks and (bottom) Breaking Bad using (a) DeepMend and (b) DeepJoin. *Fantastic Breaks* objects are shown before and after re-training, which improves results for Geometric Breaks.

erate repairs that fully restore real objects, e.g. the plates and mugs on the left of Figures 8(a) and (b). Re-training on *Fantastic Breaks* improves repairs for some objects, as shown in the middle of Figure 8, producing more holistic repairs. As shown in the right columns of Figure 8, re-training still allows generation of repairs for synthetic fractures.

To measure the accuracy of predicted repairs we use the chamfer distance (CD), given by Park et al. [41], and normal consistency (NC), given by Mescheder et al. [36]. As MendNet, DeepMend, and DeepJoin perform optimization during inference to obtain repairs, inference proceeds non-deterministically. For each approach, we pre-train Geometric Breaks or Breaking Bad, and perform re-training on *Fantastic Breaks*. For evaluation before re-training, we report quantitative metrics over 21 inference runs without re-training, totaling 21 trials. For evaluation after re-training, we report quantitative metrics over 3 re-trained models and 7 inference runs with re-training, totalling 21 trials.

Table 4 shows success rate, CD, and NC for predicted repairs using MendNet, DeepMend, and DeepJoin. As shown in the top of Table 4, re-training on the *Fantastic Breaks* dataset decreases CD and increases NC of repair parts on Geometric Breaks for DeepMend and DeepJoin. As acknowledged by the authors, MendNet struggles to restore real objects when trained on synthetic fractures. Thus re-training on 105 real samples may not benefit learning. Figures 8(a) and (b) show an improvement for the mug, plate, and cup from Geometric Breaks after re-training. Table 4 shows a decrease in CD from 0.059 to 0.051 and from 0.030 to 0.022 for DeepMend and DeepJoin respectively after re-training. Breaking Bad contains irregular objects, such as the tombstone and the model ship in Figures 8(a) and (b),

Table 4. Success rate (SR), chamfer distance (CD), and normal consistency (NC), with training on Geometric Breaks (GB, top) and on Breaking Bad (BB, bottom), before (left) and after (right) re-training on *Fantastic Breaks* (FB).

	Test	Train with GB Only			Train with GB + FB		
		SR%	CD	NC	SR%	CD	NC
MendNet	GB	99.89%	0.132	0.198	99.89%	0.205	0.185
	FB	99.68%	0.080	0.276	99.58%	0.094	0.277
DeepMend	GB	95.45%	0.063	0.607	93.65%	0.085	0.548
	FB	98.20%	0.059	0.323	99.05%	0.051	0.441
DeepJoin	GB	92.80%	0.064	0.658	95.98%	0.102	0.526
	FB	100.00%	0.030	0.426	99.89%	0.022	0.504
	Test	Train with BB Only			Train with BB + FB		
		SR%	CD	NC	SR%	CD	NC
MendNet	BB	93.23%	0.122	0.344	100.00%	0.210	0.307
	FB	95.13%	0.110	0.184	100.00%	0.116	0.302
DeepMend	BB	96.19%	0.078	0.637	92.70%	0.176	0.351
	FB	99.68%	0.065	0.213	100.00%	0.092	0.190
DeepJoin	BB	98.94%	0.042	0.688	100.00%	0.177	0.285
	FB	100.00%	0.061	0.259	100.00%	0.076	0.203

which may have more than half of the object fractured off, making it a challenging dataset for shape repair. Though re-training may improve performance on some real objects, e.g. the coaster in Figure 8, overall re-training on *Fantastic Breaks* does not improve performance, as shown on the bottom of Table 4. However, evaluation on our dataset demonstrates that, though all three repair approaches generate feasible restoration shapes for repairing purely synthetic fractures, models trained on Breaking Bad generalize less well to real fractured objects than models trained on Geometric Breaks, with a CD of 0.080 and 0.110 with MendNet, 0.059 and 0.065 with DeepMend, and of 0.030 and 0.061 with DeepJoin for Geometric Breaks and Breaking Bad respectively, as shown in the left supercolumn of Table 4.

6. Discussion

We present *Fantastic Breaks*, a novel dataset that contains full 360° 3D scans of broken objects geometrically aligned with 3D scans of their complete counterparts, with manual annotations of classes, materials, and fracture surfaces, and synthetic proxies for 3D meshes representing repair parts. The dataset continues to grow in number of physical objects and scans. The dataset is one of the first of its kind, enabling learning of the characteristics of fractured objects. *Fantastic Breaks* provides data-driven insight into fracture, overcoming the deficits of geometric approaches that make prior assumptions about the damage process that are not widely applicable, as well as the concerns of datasets based on physics simulations that are limited by current hardware in modeling real-world geometry.

An obvious limitation of our endeavor is that employing destructive processes to damage objects for the purpose of real-world data acquisition is unsustainable at a large scale. We advocate that our dataset be leveraged to learn patterns of break and internal geometric structure that are common across objects of similar materials, such as ceramics or plastics, and classes, such as mugs or cups, and to use generative approaches to conduct data-driven synthesis of breaks and internal structure given 3D models of complete objects. For instance, an interesting observation of our collection is that for reasons of cost, sustainable production, and functionality, nearly all our objects have shell rather than solid structures, whereas scans of whole objects such as bottles or statues cannot capture the internal shell structure. By exposing the internal structure, the dataset provides the opportunity to learn how to hollow out 3D models of objects, an opportunity absent from prior 3D scan datasets.

Fantastic Breaks currently contains tabletop objects acquired using desktop scanners. Larger objects may necessitate more elaborate setups, e.g., room-scale imaging systems. Future data collection can investigate the minimal number of viewpoints needed to acquire geometrically relevant understanding of internal object structure. For instance, to capture the fracture pattern of a chair leg, it may be sufficient to image the broken region from 1-2 viewpoints using a depth camera, and obtain and deform a 3D proxy from a public repository to the imaged viewpoints.

We have provided evaluations of existing approaches on automatic reconstruction of new repair parts using learning-driven approaches. The dataset is widely applicable to a range of other tasks, e.g., the broken and restoration meshes can be used to perform shape assembly cognizant of precision joins for real-world fracture boundaries. We evaluate a shape repair approach that does not require prior knowledge of the fractured region (MendNet). However, via our manual fracture surface annotations, we also plan to release the incomplete meshes devoid of the fracture surface. These incomplete meshes will benefit research in partial shape completion [16], which up until now has largely focused on synthetically generated partial shapes or on depth scans. Datasets of real-world 3D scans have contributed to significant advancements in robotic manipulation [8, 15]. *Fantastic Breaks* provides impacts in robotics research by enabling robot-driven repair, object grasp while being cognizant of fractured regions to minimize further damage, and damaged object handling for safe human-robot handover.

Acknowledgments This work is supported by the National Science Foundation (NSF) under grant IIS-2023998. We would like to thank Joseph Morrison, Nikolai Melnikov, Odin Kohler, Patrick O’Mahony, Holly Rossmann, Rosalina Delwiche, Elek Ye, and Noah Wiederhold for their assistance with object scanning.

References

- [1] Marco Attene, Michela Mortara, Michela Spagnuolo, and Bianca Falcidieno. Hierarchical convex approximation of 3d shapes for fast region selection. In *Computer graphics forum*, volume 27, pages 1323–1332. Wiley Online Library, 2008. 6
- [2] Dina Bashkirova, Mohamed Abdelfattah, Ziliang Zhu, James Akl, Fadi Alladkani, Ping Hu, Vitaly Ablavsky, Berk Calli, Sarah Adel Bargal, and Kate Saenko. Zerowaste dataset: Towards deformable object segmentation in cluttered scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21147–21157, 2022. 3
- [3] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor fusion IV: control paradigms and data structures*, volume 1611, pages 586–606. Spie, 1992. 4
- [4] Benedict J Brown, Corey Toler-Franklin, Diego Nehab, Michael Burns, David Dobkin, Andreas Vlachopoulos, Christos Doumas, Szymon Rusinkiewicz, and Tim Weyrich. A system for high-volume acquisition and matching of fresco fragments: Reassembling theran wall paintings. *ACM transactions on graphics (TOG)*, 27(3):1–9, 2008. 1, 2
- [5] Berk Calli, Arjun Singh, James Bruce, Aaron Walsman, Kurt Konolige, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. Yale-cmu-berkeley dataset for robotic manipulation research. *The International Journal of Robotics Research*, 36(3):261–268, 2017. 2, 3
- [6] Berk Calli, Aaron Walsman, Arjun Singh, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols. *IEEE Robotics and Automation Magazine*, pages 36–52, 2015. 2, 3
- [7] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5
- [8] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, et al. Dexycb: A benchmark for capturing hand grasping of objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9044–9053, 2021. 8
- [9] Yun-Chun Chen, Haoda Li, Dylan Turpin, Alec Jacobson, and Animesh Garg. Neural shape mating: Self-supervised object assembly with adversarial shape priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12724–12733, 2022. 1, 2, 3
- [10] Floyd M Chitalu, Qinghai Miao, Kartic Subr, and Taku Komura. Displacement-correlated xfem for simulating brittle fracture. In *Computer Graphics Forum*, volume 39, pages 569–583. Wiley Online Library, 2020. 2
- [11] Sungjoon Choi, Qian-Yi Zhou, Stephen Miller, and Vladlen Koltun. A large dataset of object scans. *arXiv preprint arXiv:1602.02481*, 2016. 3
- [12] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 3
- [13] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B. McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560, 2022. 2, 3
- [14] Linxu Fan, Floyd M Chitalu, and Taku Komura. Simulating brittle fracture with material points. *ACM Transactions on Graphics (TOG)*, 41(5):1–20, 2022. 2
- [15] Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11444–11453, 2020. 8
- [16] Ben Fei, Weidong Yang, Wenming Chen, Zhijun Li, Yikang Li, Tao Ma, Xing Hu, and Lipeng Ma. Comprehensive review of deep learning-based 3d point clouds completion processing and analysis. *arXiv preprint arXiv:2203.03311*, 2022. 8
- [17] Thomas Funkhouser, Hijung Shin, Corey Toler-Franklin, Antonio García Castañeda, Benedict Brown, David Dobkin, Szymon Rusinkiewicz, and Tim Weyrich. Learning how to match fresco fragments. *Journal on Computing and Cultural Heritage (JOCCH)*, 4(2):1–13, 2011. 1, 2
- [18] Robert Gregor, Danny Bauer, Ivan Sipiran, Panagiotis Perakis, and Tobias Schreck. Automatic 3d object fracturing for evaluation of partial retrieval and object restoration tasks-benchmark and application to 3d cultural heritage data. In *3DOR@ Eurographics*, pages 7–14, 2015. 1, 2, 3
- [19] Abhinav Narayan Harish, Rajendra Nagar, and Shanmuganathan Raman. Rgl-net: A recurrent graph learning framework for progressive part assembly. In *WACV*, pages 647–656, 2022. 1, 3
- [20] Renato Hermoza and Ivan Sipiran. 3d reconstruction of incomplete archaeological objects using a generative adversarial network. In *Proc. Computer Graphics International*, pages 5–11, 2018. 1, 3
- [21] Je Hyeong Hong, Young Min Kim, Koang-Chul Wi, and Jinwook Kim. Potsac: A robust axis estimator for axially symmetric pot fragments. In *ICCV Workshops*, pages 1421–1428, 2019. 1, 3
- [22] Je Hyeong Hong, Seong Jong Yoo, Muhammad Arshad Zee-shan, Young Min Kim, and Jinwook Kim. Structure-from-shards: Incremental 3d reassembly of axially symmetric pots from unordered and mixed fragment collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5443–5451, 2021. 1, 2, 3
- [23] Qi-Xing Huang, Simon Flöry, Natasha Gelfand, Michael Hofer, and Helmut Pottmann. Reassembling fractured objects by geometric matching. In *ACM SIGGRAPH*, pages 569–578, 2006. 1, 2, 3
- [24] Allison Janoch, Sergey Karayev, Yangqing Jia, Jonathan T Barron, Mario Fritz, Kate Saenko, and Trevor Darrell. A category-level 3d object dataset: Putting the kinect to work.

- In *Consumer depth cameras for computer vision*, pages 141–165. Springer, 2013. 3
- [25] Alexander Kasper, Zhixing Xue, and Rüdiger Dillmann. The kit object models database: An object model database for object recognition, localization and manipulation in service robotics. *The International Journal of Robotics Research*, 31(8):927–934, 2012. 2, 3
- [26] Kevin Lai, Liefeng Bo, Xiaofeng Ren, and Dieter Fox. Rgb-d object recognition: Features, algorithms, and a large scale benchmark. In *Consumer Depth Cameras for Computer Vision*, pages 167–192. Springer, 2013. 3
- [27] N. Lamb, S. Banerjee, and N. K. Banerjee. Automated reconstruction of smoothly joining 3d printed restorations to fix broken objects. In *Proceedings of the ACM Symposium on Computational Fabrication*, pages 1–12, 2019. 1, 2, 3, 5
- [28] N. Lamb, S. Banerjee, and N. K. Banerjee. Deepjoin: Learning a joint occupancy, signed distance, and normal field function for shape repair. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, jul 2022. 1, 2, 3, 5, 7
- [29] N. Lamb, S. Banerjee, and N. K. Banerjee. Deepmend: Learning occupancy functions to represent shape for repair. In *European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 3, 7
- [30] N. Lamb, S. Banerjee, and N. K. Banerjee. Mendnet: Restoration of fractured shapes using learned occupancy functions. *Computer Graphics Forum*, 41(5):65–78, 2022. 1, 2, 3, 7
- [31] N. Lamb, N. Wiederhold, B. Lamb, S. Banerjee, and N. K. Banerjee. Using learned visual and geometric features to retrieve complete 3d proxies for broken objects. In *Symposium on Computational Fabrication*, pages 1–15, 2021. 1, 2
- [32] Xin Li, Zhao Yin, Li Wei, Shenghua Wan, Wei Yu, and Maoqing Li. Symmetry and template guided completion of damaged skulls. *Computers & Graphics*, 35(4):885–893, 2011. 1, 3
- [33] Liu Liu, Wenqiang Xu, Haoyuan Fu, Sucheng Qian, Qiaojun Yu, Yang Han, and Cewu Lu. Akb-48: a real-world articulated object knowledge base. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14809–14818, 2022. 3
- [34] A Mandal, P Chaudhuri, and S Chaudhuri. Remeshing-free graph-based finite element method for fracture simulation. In *Computer Graphics Forum*. Wiley Online Library, 2022. 2
- [35] Pavlos Mavridis, Ivan Sipiran, Anthousis Andreadis, and Georgios Papaioannou. Object completion using k-sparse optimization. *Computer Graphics Forum*, 34(7):13–21, 2015. 1, 3
- [36] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019. 7
- [37] 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*, pages 909–918, 2019. 1
- [38] Georgios Papaioannou and Evaggelia-Aggeliki Karabassi. On the automatic assemblage of arbitrary broken solid artefacts. *Image and Vision Computing*, 21(5):401–412, 2003. 1, 3
- [39] Georgios Papaioannou, E-A Karabassi, and Theoharis Theoharis. Virtual archaeologist: Assembling the past. *IEEE Computer Graphics and Applications*, 21(2):53–59, 2001. 1, 3
- [40] Georgios Papaioannou, Tobias Schreck, Anthousis Andreadis, Pavlos Mavridis, Robert Gregor, Ivan Sipiran, and Konstantinos Vardis. From reassembly to object completion: A complete systems pipeline. *Journal on Computing and Cultural Heritage*, 10(2):1–22, 2017. 1, 3
- [41] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proc. CVPR*, pages 165–174, Piscataway, NJ, 2019. IEEE. 7
- [42] Angelia Payne, Keenan Cole, Katie Simon, Christopher Goodmaster, and Fredrick Limp. Designing the next generation virtual museum: Making 3d artifacts available for viewing and download. In *Making History Interactive: Proceedings of the 37th Annual International Conference on Computer Applications and Quantitative Methods in Archaeology (CAA)*, volume 3, pages 1–6, 2009. 2
- [43] Pedro F Proença and Pedro Simões. Taco: Trash annotations in context for litter detection. *arXiv preprint arXiv:2003.06975*, 2020. 3
- [44] Silvia Sellán, Yun-Chun Chen, Ziyi Wu, Animesh Garg, and Alec Jacobson. Breaking Bad: A dataset for geometric fracture and reassembly. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. 1, 2, 5, 6
- [45] Silvia Sellán, Jack Luong, Leticia Mattos Da Silva, Aravind Ramakrishnan, Yuchuan Yang, and Alec Jacobson. Breaking good: Fracture modes for realtime destruction. *ACM Trans. Graph.*, July 2022. 1, 2
- [46] Arjun Singh, James Sha, Karthik S Narayan, Tudor Achim, and Pieter Abbeel. Bigbird: A large-scale 3d database of object instances. In *2014 IEEE international conference on robotics and automation (ICRA)*, pages 509–516. IEEE, 2014. 2, 3
- [47] Ivan Sipiran. Completion of cultural heritage objects with rotational symmetry. In *Proceedings of the 11th Eurographics Workshop on 3D Object Retrieval*, pages 87–93, 2018. 1, 3
- [48] Ivan Sipiran, Robert Gregor, and Tobias Schreck. Approximate symmetry detection in partial 3d meshes. *Computer Graphics Forum*, 33(7):131–140, 2014. 1, 3
- [49] Joao Sousa, Ana Rebelo, and Jaime S Cardoso. Automation of waste sorting with deep learning. In *2019 XV Workshop de Visão Computacional (WVC)*, pages 43–48. IEEE, 2019. 3
- [50] Joshua Wolper, Yunuo Chen, Minchen Li, Yu Fang, Ziyin Qu, Jiecong Lu, Meggie Cheng, and Chenfanfu Jiang. Anisompm: Animating anisotropic damage mechanics: Supplemental document. *ACM Trans. Graph.*, 39, 2020. 2

- [51] Joshua Wolper, Yu Fang, Minchen Li, Jiecong Lu, Ming Gao, and Chenfanfu Jiang. Cd-mpm: continuum damage material point methods for dynamic fracture animation. *ACM Transactions on Graphics (TOG)*, 38(4):1–15, 2019. 2
- [52] Rundi Wu, Yixin Zhuang, Kai Xu, Hao Zhang, and Baoquan Chen. Pq-net: A generative part seq2seq network for 3d shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 829–838, 2020. 1, 3
- [53] Yu Xiang, Wonhui Kim, Wei Chen, Jingwei Ji, Christopher Choy, Hao Su, Roozbeh Mottaghi, Leonidas Guibas, and Silvio Savarese. Objectnet3d: A large scale database for 3d object recognition. In *European conference on computer vision*, pages 160–176. Springer, 2016. 3
- [54] Guanqi Zhan, Qingnan Fan, Kaichun Mo, Lin Shao, Baoquan Chen, Leonidas J Guibas, Hao Dong, et al. Generative 3d part assembly via dynamic graph learning. *Advances in Neural Information Processing Systems*, 33:6315–6326, 2020. 1, 3
- [55] Kang Zhang, Wuyi Yu, Mary Manhein, Warren Waggenspack, and Xin Li. 3d fragment reassembly using integrated template guidance and fracture-region matching. In *Proceedings of the IEEE international conference on computer vision*, pages 2138–2146, 2015. 1, 3
- [56] Qingnan Zhou and Alec Jacobson. Thingi10k: A dataset of 10,000 3d-printing models. *arXiv preprint arXiv:1605.04797*, 2016. 1