ACM SIGGRAPH / Eurographics Symposium on Computer Animation (2024)
M. Skouras and H. Wang (Editors)

Poster

Brittle Fracture Animation with VQ-VAE-Based Generative Method

Y. Huang[1] and T. Kanai[1]

[1]The University of Tokyo, Graduate School of Arts and Sciences, Tokyo

Abstract

We propose a novel learning-based approach for predicting fractured shapes based on collision dynamics at run-time and seamlessly integrating realistic brittle fracture animations with rigid-body simulations. Our method utilizes BEM brittle fracture simulations to create training data. We introduce generative geometric segmentation, distinct from instance and semantic segmentation, to represent 3D fracture shapes. We adopt the concept of a neural discrete representation learning framework to optimize multiple discrete fractured patterns with a continuous latent code. Additionally, we propose a novel SDF-based cage-cutting method to create fragments by cutting the original shape with the predicted fracture pattern. Our experimental results demonstrate that our approach can generate significantly more detailed brittle fractures compared to existing techniques, while reducing computational time typically when compared to traditional simulation methods at comparable resolutions.

CCS Concepts

• ***Computing methodologies*** → ***Animation; Neural networks;*** *Learning latent representations;*

1. Introduction

Brittle fracture animations bring impressive visual effects to video games, movies, and virtual reality. Most simulation methods do not focus on crack propagation but instead concentrate on determining the cutting meshes. Recently, physics-based simulation methods based on quasi-static crack propagation can generate detailed realistic fractured shapes and surfaces. However, current physics-based simulation methods based on crack propagation [HW16; FCK22], including those utilized within the film industry, suffer from computationally expensive costs.

In real-time applications like virtual reality or games, a more popular alternative involves creating a pre-fractured pattern during the modeling stage and swapping from the original shape to the fractured shape upon collision. However, the monotonous Voronoi-like shapes make it difficult to represent complex real-world fracture patterns, and these shapes are not tailored to each collision.

We introduce a novel approach for predicting brittle fracture patterns utilizing neural discrete representation learning [VV*17]. We reconceptualize the challenge of brittle fracturing as predicting a specific fracture pattern related to a BEM simulation collision condition, viewing it as a conditional 3D fracture pattern prediction.

2. Our Method

Our method consists of two main processes: 1) In the learning process, we first generate the training dataset. Then, we train a conditional generative model in the training phase of our neural discrete representation learning framework. 2) In the run-time process, we predict the fractured surfaces using the conditional generative model in the inference phase. We then synthesize the fractured surfaces and original meshes using our caged-SDF segmentation and employ the reconstructed shapes in the physics engine. Each fragment is assigned new velocities and masses derived from its state before destruction. The rigid-body simulation proceeds with these updated configurations.

Generating training data. Our framework utilizes BEM brittle fracture simulation [HW16] to generate training data that reflects collision scenarios and their resultant fracture patterns. To predict the fracture patterns, we abstract the fracturing scenarios caused by collisions and encode the scene into a deep learning-friendly dataset. This dataset pairs impulse data inputs with outputs of fractured fragments.

Training phase in learning process. We have designed a custom conditional generative model to enable the generation of fracture patterns at various 3D resolutions through a discrete representation training framework [VV*17]. The input representation includes a normalized impulse code representing the strength, direction, and position of the impulse when collision occurs. To better represent 3D fractured surfaces, we define an output 3D representation suitable for generative geometric segmentation, similar to truncated signed distance fields (TSDF), by accepting the internal unsigned distance values. In the training phase, we concatenate the input impulse code with initialized noisy code to link to specific fractured patterns and generate an embedding space following the discrete representation training framework.

DOI: 10.2312/sca.20241163

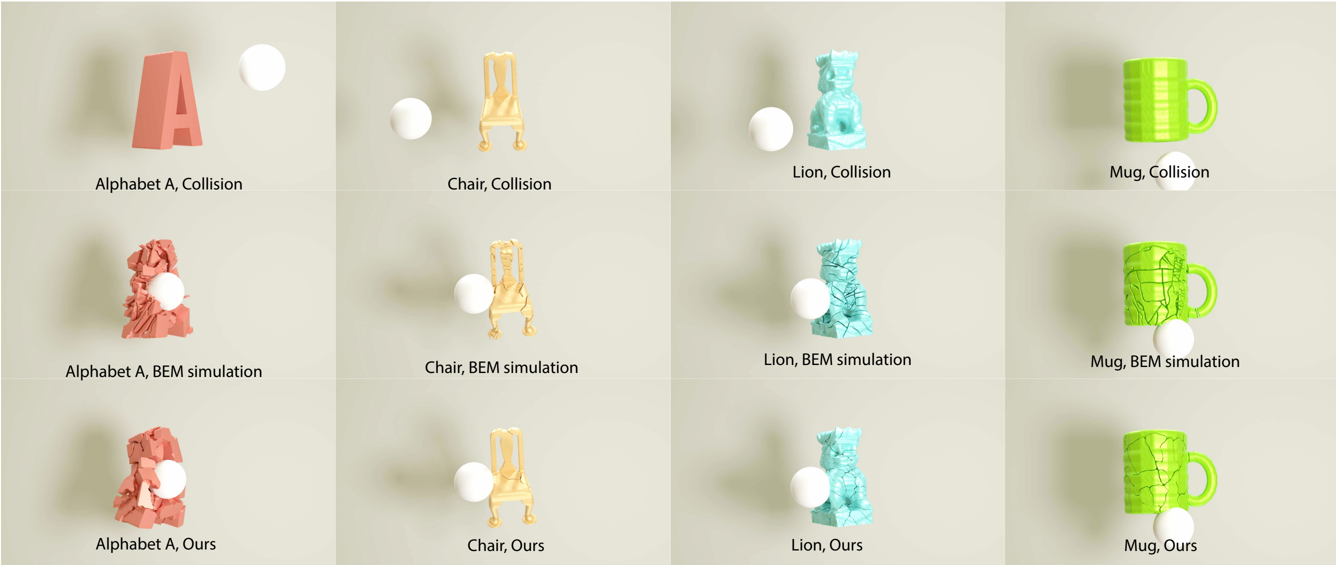

Figure 1: *Comparison between the simulation results and fractured shapes predicted by deep learning. Left to right: Alphabet A, Chair, Lion, Mug. Top to bottom row: Input collision condition, Brittle fracture simulation results, Results of our method. All collision conditions are not contained in the learning process.*

Inference Phase in run-time process. When a collision with the target shape occurs, impulse data is captured and transmitted to the conditional generative model. Similar to the training phase, we encode the normalized impulse code for the conditional generative model. In the inference phase, unlike the training phase, we generate a normal distribution random noisy code and concatenate it with the normalized impulse code. By searching for the closest element in the embedding space, we forward the conditional generative model to obtain the fractured surfaces.

Caged-SDF Segmentation. We have developed a method called *caged-SDF segmentation* to reconstruct destruction patterns while preserving the original external surface mesh. The caged-SDF segmentation method aims to generate predicted internal fractured surfaces, which are enclosed by a thin, soft-wrapping cage. This approach involves several Boolean set operations between the cage with fractured surfaces and the original mesh.

3. Results

In our experiments, we used the shape data provided by Thingi10K [ZJ16] and Objaverse 1.0 [DSS*23]. For each target shape, we conducted both the learning and run-time processes with a generative model tailored to one target shape individually. We generate the dataset for 10 hours on four PCs with Ryzen 9 5950X CPUs. We conducted over 200 frames of destruction experiments, colliding 60 times and collecting the data. As a result, we calculated 60 sets of input and output learning data for each shape, 50 sets used in the training phase, and 10 sets used in tests as shown in Figure 1. The average time cost in the training phase for one target shape is 3 hours, 46 minutes, and 14 seconds with 1500 epochs.

Figure 1 shows that our results are similar to brittle fracture simulation results in terms of destruction patterns, global complexity, and fracture surface shapes. Our deep learning-based method generated visually close results to brittle fracture simulation within an average calculation time of 12.2 seconds, compared to the average calculation time of 13.0 minutes for crack propagation-based brittle fracture simulation.

4. Conclusion and Future work

We introduced the prior application of a deep learning-based fracture system, defining the task of 3D destruction shape generation. Additionally, we develop a novel and stable SDF-based cage-cutting method that can be adapted in other works.

Future work includes further evaluating the proposed framework by improving the run-time process and developing a general generative model to predict brittle fractures across a specific category rather than focusing on individual shapes.

References

[DSS*23] DEITKE, MATT, SCHWENK, DUSTIN, SALVADOR, JORDI, et al. "Objaverse: A Universe of Annotated 3D Objects". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023, 13142–13153 2.

[FCK22] FAN, LINXU, CHITALU, FLOYD M, and KOMURA, TAKU. "Simulating Brittle Fracture with Material Points". *ACM Transactions on Graphics (TOG)* 41.5 (2022), 1–20. DOI: 10.1145/3522573 1.

[HW16] HAHN, DAVID and WOJTAN, CHRIS. "Fast Approximations for Boundary Element Based Brittle Fracture Simulation". *ACM Trans. Graph.* 35.4 (July 2016). ISSN: 0730-0301. DOI: 10.1145/2897824.2925902 1.

[VV*17] VAN DEN OORD, AARON, VINYALS, ORIOL, et al. "Neural discrete representation learning". *Advances in neural information processing systems* 30 (2017) 1.

[ZJ16] ZHOU, QINGNAN and JACOBSON, ALEC. "Thingi10K: A Dataset of 10,000 3D-Printing Models". *arXiv preprint arXiv:1605.04797* (2016) 2.