PhysGaussian: Physics-Integrated 3D Gaussians for Generative Dynamics

Jared Cambier, Shania Guo, Runfeng Li, Yiyang Zhang Original paper by: Tianyi Xie, Zeshun Zong, Yuxing Qiu, Xuan Li, Yutao Feng, Yin Yang, Chenfanfu Jiang

May 19, 2024

Abstract

For our final project in CS2240, my group decided to reimplement the paper "PhysGaussian: Physics-Integrated 3D Gaussians for Generative Dynamics" by Xie et all. In it, the authors employ the Material Point Method (MPM) to achieve motion synthesis on scenes composed of 3D gaussians. In graphics, there is often a divide between the simulating and rendering processes; for example, you might use a mesh to represent the scene, perform an intermediary step to make it simulation-ready (such as dividing it into tetrahedra), perform the simulation, and then render. In the real world, this conceptual divide between physics and "rendering" doesn't exist because forces act on the actual particles that you see. With scenes composed of 3D gaussians, we can alter their positions and covariance matrices using the material point method and render them to the screen through a process called splatting. In other words, the gaussians are used both for the simulation and rendering processeses, following the principle of "What you see is what you simulate" $((WS)^2)$. First, gaussian splats of scenes are learned using the pipeline in "3D Gaussian Splatting for Real-Time Radiance Field Rendering" by Kerbl et all. We use training images of any real-life scene, and the output is a 3D gaussian splat. To perform physics on the scene, we augment the gaussians with deformation and mechanical stress attributes. We also employ different material models, including metal, jelly, sand, foam, plastic, and non-newtonian fluid. Gravitational forces and initial particle impulses are used to test movement. Additionally, fracture is supported, and materials can collide with one another. Lastly, we propose a method for optimizing the material-specific value of Young's modulus (E), which will significantly improve the trial-and-error process of testing different values to achieve the desired visual result. If we have a ground truth video of how an object should interact with its environment, we can learn the value of E that most closely mimics its material properties (in particular, rigidity).



Left: a foam-model cake ripping apart Right: sand falling on an elastic chair