40th Anniversary of the IEEE Conference on Robotics and Automation (ICRA@40), Rotterdam, Netherlands | September 23-26, 2024

Uncertain Physics for Robot Simulation in a Game Engine

Hermann Meißenhelter[1], René Weller[1] and Gabriel Zachmann[1]

***Abstract*— Physics simulations are crucial for domains like animation and robotics, yet they are limited to deterministic simulations with precise knowledge of initial conditions. We introduce a surrogate model for simulating rigid bodies with positional uncertainty (Gaussian) and use a non-uniform sphere hierarchy for object approximation. Our model outperforms traditional sampling-based methods by several orders of magnitude in efficiency while achieving similar outcomes.**

I. Introduction

Physics simulations are pivotal in predicting complex scenarios yet often fail to mirror real-world unpredictability, resulting in a "reality gap." This gap stems from computational models' inability to capture every nuance of physics and material properties, compounded by deterministic simulators that yield identical, precise outcomes. Sampling introduces variability by allowing simulations to account for uncertain parameters, thus enhancing realism across applications like animation and robotics. For instance, software like Blender and SideFX Houdini utilize sampling to refine visuals for greater aesthetic and physical fidelity in animation. In robotics, sampling enables robots to foresee and adapt to dynamic environments, optimizing their actions for better decision-making [1].

A simulator can be divided into two modules: collision detection and simulation. The simulations receive and resolve the collisions provided by the collision detection and predict the next position of simulated objects. Physics-based animation methods can be categorized into penalty-, impulse- and constraint-based methods. The penalty-based methods use a spring-damper system to penalize penetrations [2]. It is generally usable for deformable and rigid objects [3]. On the other hand, impulse-based simulators apply collision impulses to simulate physical interactions.

Uncertainty can be propagated analytically when the function is linear. For non-linear functions, uncertainty can only be approximated. Monte Carlo sampling approximates the true distribution with an increasing number of samples. Other approximation methods include linearizing the function or using a surrogate model [4].

*This research has been (partially) supported by the German Research Foundation DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 Project-ID 329551904 "EASE - Everyday Activity Science and Engineering", University of Bremen (http://www.ease-crc.org/), subproject R03.

[1] Mathematics and Computer Science, University of Bremen, Germany
meissenhelter@cs.uni-bremen.de

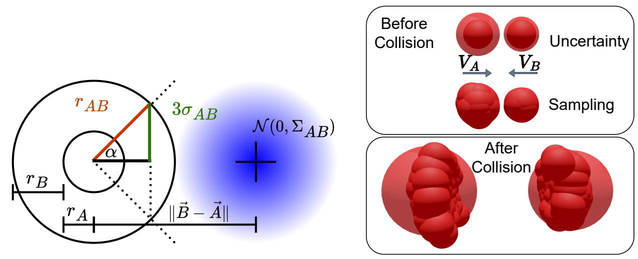

(a) Geometric Approximation (b) Directional Uncertainty

Fig. 1: Geometric approximation of the uncertainty of the contact normal (a). Dependent on the combined radii of both spheres r_{AB} and the convoluted uncertainty of both spheres σ_{AB} (right triangle). Approximation applied for two colliding spheres and compared to sampling (1k spheres) (b). Three-sigma-hull for positional uncertainty in light red.

II. Probabilistic Rigid Body Simulation

A. Collision Detection

We approximate polygonal rigid bodies using non-uniform spheres, leveraging their low computational cost for intersection tests and independence from polygonal resolution. We employ an Apollonian sphere packing algorithm—adapted from [5] for arbitrary 3D objects—that uses a greedy approach favoring larger spheres to achieve space-filling packing without overlaps within the 3D object. Subsequently, each sphere packing is organized into a bounding volume hierarchy, enabling rapid intersection detection for complex 3D objects.

B. Simulation

The state vector of a rigid body is defined within $\mathrm{SE}(3)$. We currently model the state's uncertainty only in linear motion, and the uncertainty is constrained to be isotropic.

When a collision occurs, we resolve it with either our penalty or the impulse approach. The penalty method models a spring, where the force scales with the penetration volume. We go through each contact point and accumulate the penalty force. The impulse approach models the physical interaction using collision impulses. We sequentially compute and apply the impulse for each contact point, taking into account the velocities, masses, and normals involved in the contact. Our surrogate model approximates the deviation from the contact normal as a cone angle geometrically (see Fig. 1). This angle is then used to compute the tangential component of velocity or force. However, we combine the uncertainty in direction by accumulating each contact's angle and clamping it to 90 degrees as an upper bound. The kinematic uncertainty propagation is physically motivated. We use the symplectic Euler, known for its stability and simplicity. Additionally, we clamp the uncertainty in velocity σ_v by the mean velocity v magnitude. Clamping resulted in more reasonable growth of uncertainty in general. As a second method, we computed the first-order Jacobi of the non-linear collision resolution function and propagated the uncertainty similarly, as the Extended Kalman Filter does. Here, we use auto-differentiation.

III. Evaluation

In the following, we evaluate our approach in terms of performance and quality. We have simulated a collision between three 3D models—Armadillo, Bunny, and Cup. A pair of objects approaches with an uncertain initial position and certain velocity. There is no gravity or other collision, i.e., with a floor involved. However, a linear damping for the velocity is applied. We simulated 5 s with a $\Delta t = 0.01$, s. We varied the number of spheres, method, and positional sigma during the experiments. The models are scaled to fit inside a unit-sphere so the objects are similarly sized, and an equal absolute sigma is credible. We examine scenarios with 10^n spheres and samples, where $n = 1, 2, 3, 4, 5$, corresponding up to 100k spheres and samples. The maximum sigma is 10% of the max extent (0.1) and uniformly divided. This means each object has the same absolute sigma. We use the sampling approach for comparison, where 100k samples act as a ground truth.

A. Performance Analysis

We first grouped the results by object type and number of spheres to evaluate the performance. Then, we accumulated the computation time of each frame for each sigma. Fig. 2 shows the performance. In that plot, we can see that the sampling approach scales linearly in terms of performance for each number of samples. On the other hand, it stays constant for each sigma and is therefore independent of it. However, our approach is up 3-4 orders of magnitude faster.

B. Similarity Analysis

Our final distribution is isotropic Gaussian, and the sampling distribution might not be Gaussian. However, we fit an isotropic Gaussian for the sampling distribution, as it is the best solution achievable and comparable to our approach. To compare two multivariate Gaussians for similarity, we use the Hellinger distance. Its advantage is that it is metric, and the resulting distance d is in the interval $0 \leq d \leq 1$. Where zero distance indicates similar distribution and distance of one indicates no similarity. In Fig. 3 we can see the average similarity for each sigma where 100k samples are the reference. With more samples, the similarity to the ground truth increases, but not linearly. Our surrogate model achieves moderate similarity and performs better than the linearization approach. For the linearization approach, we compute and compare with its largest eigenvalue of the covariance since it results in a non-isotropic distribution.

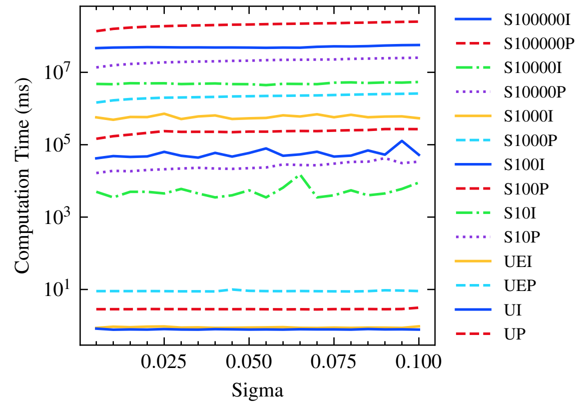

Fig. 2: Computation time over sigma for different methods. Armadillo with 10k spheres. Our approach: below 10 ms. The following abbreviations are used : (S, Sampling), (U, UncertainPhysics), (UE, Linearization) (P, Penalty method), (I, Impulse method). After sampling (S), the number indicates the number of samples.

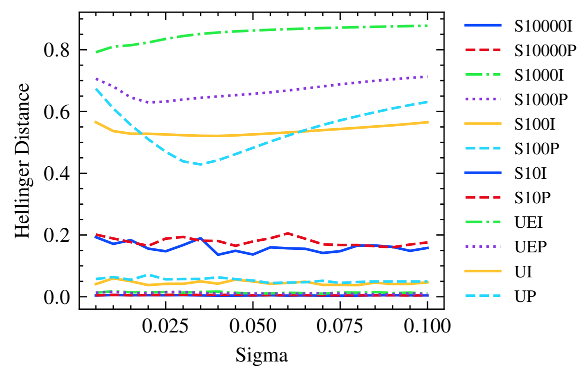

Fig. 3: Average similarity to the 100k sampling distribution over sigma for different objects and number of spheres. Ten samples have ≈ 0.2 distance. Same abbrv. as in Fig. 2.

IV. Conclusion and Future Work

This study has presented a novel approach to addressing computationally intensive sampling-based physics simulations with uncertainty. Our approach helps to reduce the total number of samples needed (i.e., in an animation or robotic context). Since we can achieve an enormous speedup, there is room for further quality improvement by experiments with more computationally intensive approaches, i.e., non-isotropic Gaussian, including rotational uncertainty and multi-modal distributions.

References

[1] F. K. Kenghagho, M. Neumann, P. Mania, T. Tan, F. A. Siddiky, R. Weller, G. Zachmann, and M. Beetz, "Naivphys4rp-towards human-like robot perception "physical reasoning based on embodied probabilistic simulation"," in *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*. IEEE, 2022, pp. 815–822.

[2] M. Tang, D. Manocha, M. A. Otaduy, and R. Tong, "Continuous penalty forces," *ACM Transactions on Graphics (TOG)*, vol. 31, no. 4, pp. 1–9, 2012.

[3] J. Jansson and J. S. Vergeest, "A general mechanics model for systems of deformable solids," in *Proceedings of TMCE*, 2000.

[4] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, “A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking,” *IEEE Transactions on signal processing*, vol. 50, no. 2, pp. 174–188, 2002.

[5] R. Weller and G. Zachmann, “Protosphere: A gpu-assisted prototype guided sphere packing algorithm for arbitrary objects,” in *ACM SIGGRAPH ASIA 2010 Sketches*, ser. SA ’10. New York, NY, USA: Association for Computing Machinery, 2010. [Online]. Available: https://doi.org/10.1145/1899950.1899958