ISSN 2464-4617 (print)
ISSN 2464-4625 (online) Computer Science Research Notes - CSRN http://www.wscg.eu

Exploring Approaches for Rigid Body Dynamics with Uncertainty

Hermann Meißenhelter, Franklin Kenghagho Kenfack, René Weller, Gabriel Zachmann
University of Bremen, Germany
{meissenhelter, fkenghag, weller, zach}@cs.uni-bremen.de

Abstract

Physics simulations are crucial for domains such as animation and robotics, yet they are limited to deterministic simulations that require precise knowledge of initial conditions. Uncertainty often requires computationally intensive sampling. This paper addresses this challenge by incorporating position and orientation uncertainty into a rigid-body simulation. We formulate and compare four uncertainty propagation approaches: linearization, a surrogate model, persistent sigma points, and local sampling. Our methods build on a penalty-based simulator with a non-uniform sphere hierarchy for geometry approximation and collision detection. We evaluate these methods against Monte Carlo sampling in collision-free and collision scenarios using the Wasserstein distance, and compare their storage requirements. Across our test cases, the proposed methods are about an order of magnitude faster than sampling-based baselines while preserving competitive accuracy. The surrogate model is the fastest, whereas persistent sigma points achieve the best overall accuracy.

Keywords
Physics-based animation, Uncertainty, Collision, Gaussian, Sampling, Monte Carlo

1 INTRODUCTION

Physics simulations are central to predicting complex scenarios but often fail to capture 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, thereby enhancing realism across applications such as animation and robotics. For instance, software such as Blender and SideFX Houdini uses 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 [KNM+22]. Similarly, recent research in cognition posits that people possess an Intuitive Physics Engine (IPE) [SHS+24], analogous to a game engine, that simulates the world. The IPE takes the current state of the world as input and yields predictions about future states. In studies, future simulations are sampled since game engines are not probabilistic (see Fig.1).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

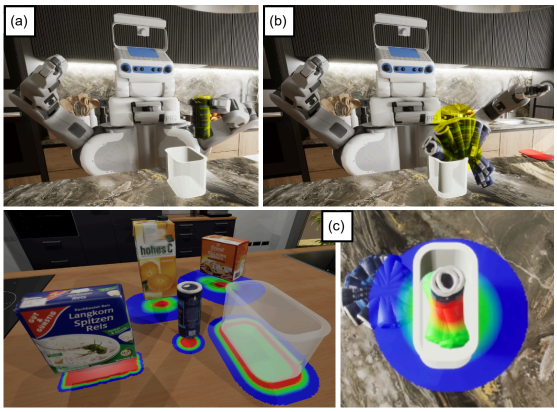

Figure 1: Teaser: (a) A robot drops a salt box into a bin under positional uncertainty. (b) A sampling-based mental simulation predicts possible outcomes using many rollouts, which is computationally expensive. (c) Our method incorporates Gaussian uncertainty directly into the simulation and predicts the resulting outcome distribution much faster than sampling.

To the best of our knowledge, prior work has not addressed uncertainty propagation for rigid-body collisions in this form. Uncertainty propagation can be performed analytically in cases where the transformation function is linear, forming the mathematical basis for the Kalman filter. For non-linear functions, however, uncertainty must be approximated in general. Monte Carlo sampling, for instance, increasingly approximates the actual distribution as the number of samples grows. Other approximation methods include function linearization or the use of a surrogate model [AMGC02].

This paper studies probabilistic rigid-body simulation under Gaussian uncertainty in position and orientation, where collisions make the dynamics strongly nonlinear and challenging for standard uncertainty-propagation techniques. We focus on penalty-based collision dynamics with non-uniform sphere hierarchies and investigate how such uncertainty can be propagated more efficiently than by repeated Monte Carlo sampling.

Our contributions are:

- A collision-aware linearization method that propagates joint cross-body covariance through the full rigid-body simulator.
- A square-root surrogate model that approximates the growth of contact-induced uncertainty geometrically by converting uncertain contact configurations into velocity and angular-velocity uncertainty.
- A persistent sigma-point method that propagates sigma points through the full simulator without regenerating them after each time step, which is empirically more stable in collision settings.
- A local sampling method that estimates collision-induced force and torque covariance from perturbations around the current contact configuration, avoiding full trajectory sampling.
- A benchmark of all four approaches against Monte Carlo baselines in collision-free and collision scenarios, comparing runtime, storage, and distributional accuracy using the Wasserstein distance.

2 RELATED WORK

Physics-based animation. Simulating rigid-body motion requires careful consideration of various parameters, such as the initial positions, velocities, and stiffness of each object involved. It is also essential to consider the surface normals at each collision. To better model the inherent inhomogeneity of everyday objects, one solution is to incorporate parameter variations in simulations, resulting in a more authentic simulation [BHW96]. When a particular goal needs to be achieved, gradient descent-based methods can be applied. However, collisions result in a challenge for gradient methods because the motion changes abruptly. This problem increases with more collisions, which can make forms of sampling necessary. Other work focuses on making sampling-based approaches more interactively explorable [TJ07] and combining different simulators [GJ22]. 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 [TPBF87, TMOT12]. It is generally usable for deformable and rigid objects [JV00]. On the other hand, impulse-based simulators apply collision impulses to simulate physical interactions. Constraint-based approaches formulate collision and other mechanical constraints (e.g., joint constraints, non-penetration constraints) within a system of equations that must remain satisfied at each simulation step [Wit97]. However, due to the complexity and nonlinearity of these systems, closed-form analytic solutions are generally not available. Instead, iterative solvers are used to approximate solutions and ensure that the constraints are maintained throughout the simulation.

Dynamics under uncertainty. Uncertainty-aware simulation is closely related to recursive state estimation, where the Kalman filter provides the canonical linear-Gaussian baseline for propagating means and covariances through system dynamics [Kal60]. For nonlinear but still smooth models, Kalman-based methods have typically relied either on local linearization, as in the extended Kalman filter, or on sigma-point propagation, as in the unscented Kalman filter [JU97]. Sampling-based approaches such as particle filters relax Gaussian assumptions and can represent multimodal posteriors, making them favorable when uncertainty grows strongly under nonlinear dynamics [GSS93]. However, these classical estimators are usually formulated for smooth state-space models and, by themselves, do not account for event-driven changes in the flow. This becomes critical in rigid-body simulation, where impacts and contact changes induce discontinuous velocity updates and strong sensitivity to event timing. Additionally, as noted in [KPCJ21], the number of particles needed might grow exponentially. In the hybrid-systems literature, the saltation matrix provides the first-order correction needed to propagate perturbations across such discontinuities rather than only through the continuous dynamics [KJPZJ24]. In contrast, we use a smooth penalty-based method. Recent work extending Kalman-based estimation to hybrid systems explicitly notes that comparatively little work has addressed how probability distributions propagate through hybrid transitions [KPCJ21]. A significant difference from prior work is that we have no external (sensor) measurements to fuse into the uncertainty, so our uncertainty typically grows over time.

3 PROBABILISTIC RIGID-BODY SIMULATION

In the following, we first define the rigid body state, then describe the underlying physics model, and finally present our methods for simulation under uncertainty.

3.1 State and Dynamics

Each rigid body is described by the mean state

$$\mathbf{x} = (\mathbf{p}, \mathbf{v}, \mathbf{R}, \omega), \qquad \mathbf{P} \in \mathbb{R}^{12\times 12}, \tag{1}$$

where $\mathbf{p} \in \mathbb{R}^3$ is position, $\mathbf{v} \in \mathbb{R}^3$ is linear velocity, $\mathbf{R}$ is orientation, and $\omega \in \mathbb{R}^3$ is angular velocity expressed in the world frame. The covariance $\mathbf{P}$ is defined over the error state

$$\delta\mathbf{x} = \begin{bmatrix} \delta\mathbf{p}^\top & \delta\mathbf{v}^\top & \delta\theta^\top & \delta\omega^\top \end{bmatrix}^\top, \tag{2}$$

where $\delta\theta \in \mathbb{R}^3$ denotes a small orientation error expressed in the body frame.

The representation of uncertainty regarding a position using Gaussian distributions is relatively straightforward, as the position in three-dimensional space is directly expressible as a vector of real numbers. These vectors are well-suited to Gaussian statistics, where the mean and covariance matrix can effectively describe the dispersion of positions around a central point. However, representing uncertainty in rotation is more complex. Rotation in three dimensions typically involves parameters that must respect certain non-linear constraints, such as unit norm. This non-Euclidean nature of the rotation space complicates the direct application of Gaussian distributions, as the Gaussian model assumes a linear space. One approach involves linearizing [AMGC02] a parametrization, such as a quaternion, using a Jacobian matrix, which is also necessary for normalization. This process results in an effective 4x4 covariance matrix, which is difficult to interpret and visualize. However, a more recent method [LWMC12] models the uncertainty within the tangential space of $SO(3)$ (set of all rotations), taking advantage of the Euclidean properties of this space. We can transform elements from the Lie algebra, denoted as $\omega \in \mathfrak{so}(3)$, to the corresponding Lie group using the exponential map, where $\exp(\omega) \in SO(3)$. Conversely, elements from the Lie group, represented as $m \in SO(3)$, can be converted back to the Lie algebra using the logarithmic map, $\log(m) \in \mathfrak{so}(3)$. The Lie algebra of a Lie group can be understood as the tangent space at the identity element of the group. Therefore, rotations about a mean rotation $\bar{\theta}$ can be sampled as follows: $\exp(\varepsilon)\bar{\theta}$, where $\varepsilon \sim \mathcal{N}(0, \Sigma_\theta)$ represents small rotations in the Lie algebra.

The simulator consists of symplectic Euler integration followed by discrete collision detection and resolution, and can be viewed as a discrete-time transition function

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t). \tag{3}$$

For simplicity we use linear damping of velocity and angular velocity.

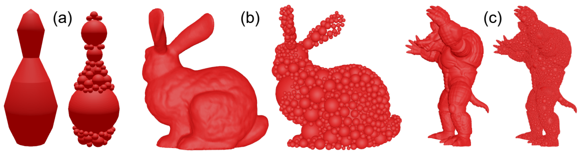

Figure 2: We use sphere packings, which approximate objects for collision detection. (a) Pin with 100 spheres, (b) Bunny with 1,000 spheres, and (c) Armadillo with 100,000 spheres.

3.2 Collision Detection and Response

We use non-uniform spheres to approximate polygonal rigid bodies. An intersection test between spheres has a negligible computational cost, and we gain independence of the polygonal resolution. An Apollonian sphere packing is a fractal structure based on rules. That idea was extended in [WZ10] to arbitrary 3D objects and is the algorithm we use to create a sphere packing. It is a greedy algorithm that prefers larger spheres. The result is a potentially space-filling sphere packing (see Figure 2). The spheres do not overlap and are entirely inside the 3D object. In the next step, we compute for each sphere packing a bounding volume hierarchy as in [WZ09], where it is called *Inner Sphere Trees*, that can detect intersection at haptic rates for complex 3D objects.

When two rigid bodies intersect, we collect all overlapping pairs of spheres. Each overlapping sphere pair defines one contact. For a contact k, the penalty force is computed as

$$\mathbf{f}_k = \kappa V_k \mathbf{n}_k, \tag{4}$$

where κ is the stiffness, V_k is the intersection volume of the two spheres, and $\mathbf{n}_k$ is the contact normal.

The resulting torque on rigid body i is

$$\tau_{i,k} = (\mathbf{c}_k - \mathbf{p}_i) \times \mathbf{f}_k, \tag{5}$$

where $\mathbf{c}_k$ is the contact point and $\mathbf{p}_i$ is the center of mass of the rigid body. The contact forces and torques are summed over all contacts and then integrated explicitly:

$$\mathbf{v}_{t+1} = \mathbf{v}_t + \Delta t\, m_i^{-1}\mathbf{f}_i, \qquad \omega_{t+1} = \omega_t + \Delta t\, \mathbf{I}_i^{-1}\tau_i. \tag{6}$$

4 OUR UNCERTAINTY PROPAGATION METHODS

Table 1 provides a high-level comparison of the proposed uncertainty propagation methods, that we introduce in the next sections. The main differences lie in

Method	Mean update	Uncertainty	Cross-body corr.	Sep. int./coll.	Collision treatment
Linearization	nominal	full covariance	yes	no	linearized
Surrogate	nominal	square-root, per body	no	yes	surrogate
Persistent sigma points	estimated	full covariance	yes	no	simulated
Local sampling	nominal	covariance, per body	no	yes	sampled

Table 1: Overview of the four uncertainty propagation methods, comparing mean propagation, uncertainty representation, cross-body correlations, separation of integration and collision handling, and collision treatment.

the uncertainty representation, the treatment of cross-correlations between rigid bodies, and whether integration and collision handling are modeled jointly or separately. Additionally, beside the sigma-point approach, our methods model the penalty force as a constant at each contact point. Omitting this assumption introduces additional uncertainty, which, in practice, results in significant uncertainty along the contact normal for each method, but the objects typically repel.

4.1 Linearization

Our first approach linearizes the non-linear transition function in Eq. (3). Given a nominal state $\mathbf{x}_t$ and a small perturbation $\delta\mathbf{x}_t$, we define

$$\delta\mathbf{x}_{t+1} = g_t(\delta\mathbf{x}_t) = f(\mathbf{x}_t \boxplus \delta\mathbf{x}_t) \boxminus f(\mathbf{x}_t). \qquad (7)$$

Here, following the $\boxplus/\boxminus$ formulation of Hertzberg et al. [HWFS13], $\boxplus$ applies a perturbation to the nominal state and $\boxminus$ maps the difference between two states back to the local error-state representation. Linearizing around $\delta\mathbf{x}_t = \mathbf{0}$ yields

$$\delta\mathbf{x}_{t+1} \approx \mathbf{F}_t\, \delta\mathbf{x}_t, \qquad \mathbf{F}_t = \left.\frac{\partial g_t}{\partial \delta\mathbf{x}_t}\right|_{\delta\mathbf{x}_t=\mathbf{0}}. \qquad (8)$$

In practice, we have computed $\mathbf{F}_t$ numerically by finite differences and also using automatic differentiation, with only a slight difference in runtime but similar accuracy. For two rigid bodies, we stack both local error states into a joint state $\delta\mathbf{x}_t = [(\delta\mathbf{x}_t^{(1)})^\top\, (\delta\mathbf{x}_t^{(2)})^\top]^\top \in \mathbb{R}^{24}$ and propagate the full covariance $\mathbf{P}_t \in \mathbb{R}^{24\times24}$, whose off-diagonal blocks model cross-correlations between the bodies. The covariance is propagated according to

$$\mathbf{P}_{t+1} = \mathbf{F}_t\mathbf{P}_t\mathbf{F}_t^\top. \qquad (9)$$

To improve numerical stability, we apply two modifications to the physics function. First, the contact normal from Eq. (4) is regularized as

$$\tilde{\mathbf{n}} = \frac{\mathbf{n}}{\|\mathbf{n}\| + r_1 + r_2}, \qquad (10)$$

since direct normalization was found to cause large covariance growth. Second, instead of resolving all detected contacts, we retain only the contact with the largest overlap volume for each rigid-body pair, which significantly reduces covariance inflation. This dominant-contact restriction is used only in the linearization approach and the other approaches retain and process all detected contacts.

4.2 Surrogate Model

The key idea of our surrogate model is to geometrically approximate the deviation from the contact normal $\mathbf{n}$ as a cone angle, as depicted in Figure 3. This angle is subsequently used to calculate the tangential force component at a contact point. The resulting tangential force is then employed to scale the uncertainty perpendicular to the contact normal. Our surrogate model stores uncertainty in square-root form. This helps to model linear growth without the needed correlations in variance space. Instead of storing the covariance matrix $\mathbf{P}$ directly, we store a matrix $\mathbf{S}$ such that

$$\mathbf{P} = \mathbf{S}\mathbf{S}^\top. \qquad (11)$$

Given an initial covariance $\mathbf{P}_0$, we compute its symmetric square-root factor by eigendecomposition

$$\mathbf{P}_0 = \mathbf{Q}\Lambda\mathbf{Q}^\top, \qquad \mathbf{S}_0 = \mathbf{Q}\Lambda^{1/2}\mathbf{Q}^\top. \qquad (12)$$

For this model, we propagate each rigid body independently and do not represent cross-correlations between different rigid bodies. For the integration step, we use the linearized single-body transition

$$\delta\mathbf{x}_{t+1} \approx \mathbf{A}_t\, \delta\mathbf{x}_t, \qquad (13)$$

with

$$\mathbf{A}_t = \begin{bmatrix} \mathbf{I} & \Delta t\,\mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \gamma\mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} & \Delta t\,\mathbf{R}_t^\top \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \gamma\mathbf{I} \end{bmatrix}. \qquad (14)$$

The square-root covariance is then propagated as

$$\mathbf{S}_{t+1} = \mathbf{A}_t\mathbf{S}_t. \qquad (15)$$

For collisions, uncertainty is injected directly in square-root space. This yields approximately linear growth in standard-deviation space and was found to be more stable than updating the full covariance directly.

For each contact between rigid bodies i and j, we first estimate the uncertainty of the two contact points. Let $\Sigma_p^{(i)}$ and $\Sigma_\theta^{(i)}$ denote the position and rotation covariance blocks of body i, and let $\mathbf{R}_i$ be its nominal orientation. With local contact offset $\mathbf{r}_{i,\mathrm{loc}}$, the contact-point covariance is approximated by

$$\Sigma_c^{(i)} = \Sigma_p^{(i)} + \mathbf{J}_i\Sigma_\theta^{(i)}\mathbf{J}_i^\top, \qquad \mathbf{J}_i = -\mathbf{R}_i[\mathbf{r}_{i,\mathrm{loc}}]_\times. \qquad (16)$$

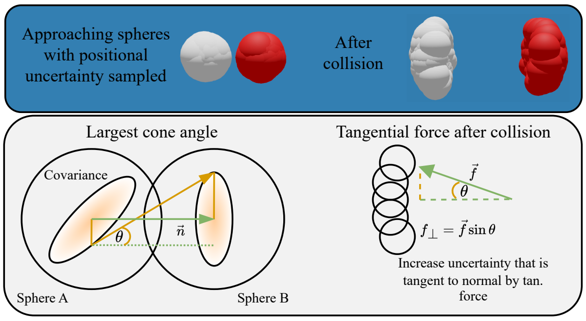

Figure 3: Illustration of the surrogate model. Top: sampled sphere configurations after collision, showing the spread induced by positional uncertainty. Bottom: schematic of our surrogate idea, where uncertainty at contact is converted into tangential uncertainty growth using a cone-angle approximation.

Analogously, we obtain $\Sigma_c^{(j)}$ for body j, and define the relative contact covariance

$$\Sigma_{\text{rel}} = \Sigma_c^{(i)} + \Sigma_c^{(j)}. \tag{17}$$

Let μ_i and μ_j be the nominal world-space contact points,

$$\mathbf{n} = \frac{\mu_j - \mu_i}{\|\mu_j - \mu_i\|}, \qquad d = \|\mu_j - \mu_i\|, \tag{18}$$

and let $\mathbf{U} = [\mathbf{u}_1\ \mathbf{u}_2]$ be an orthonormal basis of the tangent plane orthogonal to $\mathbf{n}$. Projecting the relative covariance into the tangent plane gives

$$\Sigma_T = \mathbf{U}^{\top} \Sigma_{\text{rel}} \mathbf{U}. \tag{19}$$

Let λ_1, λ_2 be the eigenvalues of Σ_T with corresponding principal directions $\mathbf{u}_1, \mathbf{u}_2$. The associated worst-case angular deviations are approximated by

$$\alpha_k = \tan^{-1}\left(\frac{k_\sigma \sqrt{\lambda_k}}{d}\right), \qquad k \in \{1,2\}, \tag{20}$$

where k_σ is the chosen confidence multiplier.

Using the penetration volume V_k and stiffness κ, we define the tangential square-root increment

$$\mathbf{V} = \sigma_1 \mathbf{u}_1 \mathbf{u}_1^{\top} + \sigma_2 \mathbf{u}_2 \mathbf{u}_2^{\top}, \qquad \sigma_k = \frac{\kappa V_k \sin(\alpha_k)}{3}. \tag{21}$$

This term is added directly to the velocity block of each rigid body,

$$\mathbf{S}_{vv}^{(i)} \leftarrow \mathbf{S}_{vv}^{(i)} + \Delta t\, m_i^{-1} \mathbf{V}, \qquad \mathbf{S}_{vv}^{(j)} \leftarrow \mathbf{S}_{vv}^{(j)} + \Delta t\, m_j^{-1} \mathbf{V}, \tag{22}$$

where m_i and m_j are the body masses.

Angular-velocity uncertainty is updated analogously from the lever arm $\mathbf{r}_i = \mu_i - \mathbf{p}_i$ and the world-frame inverse inertia tensor $\mathbf{I}_i^{-1}$:

$$\mathbf{S}_{\omega\omega}^{(i)} \leftarrow \mathbf{S}_{\omega\omega}^{(i)} + \Delta t\, \mathbf{I}_i^{-1} [\mathbf{r}_i]_\times \frac{\mathbf{V}}{3}. \tag{23}$$

An analogous update is applied to body j. This update is a approximation designed for stable uncertainty growth during contact by injecting only the tangential part of a local contact penalty scale. The angle α_k estimates how strongly uncertainty in the contact configuration may tilt the contact direction away from the nominal normal. We convert this angular deviation into a tangential uncertainty increment by multiplying the scale $\kappa V_k/3$ with $\sin(\alpha_k)$ (see Eq. 21).

Because $\alpha_k = \tan^{-1}(k_\sigma \sqrt{\lambda_k}/d)$, we have $\alpha_k \in [0, \pi/2)$ and therefore $\sin(\alpha_k) \in [0,1)$. Hence, for fixed contact stiffness and penetration volume, each tangential increment is bounded by

$$0 \leq \sigma_k < \frac{\kappa V_k}{3}. \tag{24}$$

The surrogate thus increases velocity uncertainty proportionally to the estimated deviation of the contact normal, while preventing unbounded uncertainty injection from a single contact update.

4.3 Persistent Sigma-Point Method

Our next approach is a sigma-point method inspired by the unscented transform. For two rigid bodies, we use the joint local state $\delta\mathbf{x}_t \in \mathbb{R}^{24}$, resulting in $2n+1 = 49$ sigma points. At initialization, the sigma-point offsets are generated from the initial mean $\bar{\mathbf{x}}_0$ and covariance $\mathbf{P}_0$ as

$$\delta\mathbf{x}_0^{(0)} = \mathbf{0}, \qquad \delta\mathbf{x}_0^{(j)} = \begin{cases} \mathbf{a}_j, & j = 1,\ldots,n, \\ -\mathbf{a}_{j-n}, & j = n+1,\ldots,2n, \end{cases} \tag{25}$$

where $\mathbf{A} = [\mathbf{a}_1, \ldots, \mathbf{a}_n]$ satisfies

$$\mathbf{A}\mathbf{A}^{\top} = (n+\lambda)\mathbf{P}_0. \tag{26}$$

The corresponding state sigma points are

$$\mathbf{x}_0^{(j)} = \bar{\mathbf{x}}_0 \boxplus \delta\mathbf{x}_0^{(j)}. \tag{27}$$

Unlike the standard unscented transform, we do not regenerate sigma points after each time step. Instead, the sigma points are propagated persistently through the full simulator,

$$\mathbf{x}_{t+1}^{(j)} = f(\mathbf{x}_t^{(j)}), \qquad j = 0,\ldots,2n. \tag{28}$$

Empirically, regenerating sigma points during collisions introduced large errors, whereas persistent sigma points were significantly more stable.

After propagation, the mean is recovered from the transformed sigma points using uniform weights $w_j = 1/(2n+1)$. Consistently, we set $\lambda = \frac{1}{2}$ in Eq. 26, so that the symmetric sigma point construction exactly recovers the initial covariance. Translational quantities are averaged linearly, while rotations are averaged per rigid body using an iterative Karcher mean. Rotational deviations are then mapped to $\mathbb{R}^3$ with the quaternion log map, and the joint covariance is reconstructed as

$$\mathbf{P}_{t+1} = \sum_{j=0}^{2n} w_j \, \delta\mathbf{x}_{t+1}^{(j)} \delta\mathbf{x}_{t+1}^{(j)\top}. \tag{29}$$

Because the covariance is reconstructed from the full 24-dimensional joint state, cross-correlations between the two rigid bodies are captured implicitly.

4.4 Local Sampling

Our final approach also propagates each rigid body independently and does not model cross-correlations between different rigid bodies. As in the square-root surrogate model, the collision-free integration step is propagated with the linearized single-body transition

$$\mathbf{P}_{t+1} = \mathbf{A}_t \mathbf{P}_t \mathbf{A}_t^\top. \tag{30}$$

To account for uncertainty introduced by collisions, we estimate a local covariance increment by sampling perturbations around the current contact configuration. For a contact between rigid bodies i and j, we draw samples

$$\delta\mathbf{p}_i^{(s)} \sim \mathscr{N}(\mathbf{0}, \mathbf{P}_{pp}^{(i)}), \qquad \delta\theta_i^{(s)} \sim \mathscr{N}(\mathbf{0}, \mathbf{P}_{\theta\theta}^{(i)}), \tag{31}$$

and analogously for body j, where $\mathbf{P}_{pp}$ and $\mathbf{P}_{\theta\theta}$ denote the position and rotation covariance blocks.

For each sample, the penalty-based collision response is evaluated using the perturbed contact geometry, resulting in sampled contact forces $\mathbf{f}^{(s)}$ and torques $\tau_i^{(s)}$, $\tau_j^{(s)}$. Using the nominal contact response $(\mathbf{f}^{(0)}, \tau_i^{(0)}, \tau_j^{(0)})$ as reference, we estimate the corresponding covariance increments as

$$\Sigma_f = \frac{1}{N} \sum_{s=1}^{N} \left(\mathbf{f}^{(s)} - \mathbf{f}^{(0)}\right) \left(\mathbf{f}^{(s)} - \mathbf{f}^{(0)}\right)^\top, \tag{32}$$

$$\Sigma_{\tau_i} = \frac{1}{N} \sum_{s=1}^{N} \left(\tau_i^{(s)} - \tau_i^{(0)}\right) \left(\tau_i^{(s)} - \tau_i^{(0)}\right)^\top, \tag{33}$$

and analogously for Σ_{τ_j}.

These quantities are then mapped to velocity and angular-velocity covariance updates:

$$\mathbf{P}_{vv}^{(i)} \leftarrow \mathbf{P}_{vv}^{(i)} + (\Delta t \, m_i^{-1})^2 \Sigma_f, \tag{34}$$

$$\mathbf{P}_{\omega\omega}^{(i)} \leftarrow \mathbf{P}_{\omega\omega}^{(i)} + \Delta t^2 \mathbf{I}_i^{-1} \Sigma_{\tau_i} \mathbf{I}_i^{-\top}, \tag{35}$$

with an analogous update for body j. Hence, collision uncertainty is approximated locally by sampling the variability of the contact force and torque induced by the current state uncertainty.

5 EVALUATION

In this section, we compare our predicted distribution with the Monte Carlo distribution, which we treat as the ground truth. Our final distribution is modeled as a non-isotropic Gaussian, whereas the Monte Carlo sampling distribution is not necessarily Gaussian. To quantify the discrepancy between these two distributions, we use the Wasserstein distance. A key advantage of this metric is that it does not rely on restrictive assumptions about the underlying distributions. Moreover, it has a natural interpretation as the average distance that must be transported to transform one distribution into the other, or, equivalently, the amount of work required for this transformation.

The main drawback of the Wasserstein distance is its high computational cost. Therefore, we compute a regularized W_1 distance using the Python Optimal Transport library [FCG$^+$21, FVCC$^+$24]. Since our prediction is parameterized as a Gaussian distribution, we generate samples from it to match the ground-truth distribution's sample count. We then compute the Wasserstein distance between the resulting two discrete distributions. For translational components, we use the Euclidean norm as the ground cost, while for rotational components, we use the geodesic distance measured in degrees. Our computational results were generated using an Intel® Core™ i7-12700K CPU, and we are taking the last frame for comparison, which has the highest dissimilarity.

5.1 Collision-Free Rigid-Body Motion

To validate the implementation, we first evaluate all methods in a collision-free setting; the results are shown in Table 2. In this case, uncertainty propagation is close to linear. Therefore, for uncertainty in linear velocity, the propagated distribution should theoretically be recovered without error, since Gaussian distributions are closed under linear transformations. This is reflected in the very small positional errors of the analytic methods.

For uncertainty in angular velocity, a small error is still expected due to the non-Euclidean nature of rotations and the associated approximation in the rotational representation. Nevertheless, these errors remain relatively small compared to the collision scenarios considered later. Overall, this experiment serves primarily as a sanity check and confirms that the proposed methods behave as expected in the absence of contact.

5.2 Rigid-Body Collisions

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 Pin. A pair of objects approaches with an uncertain initial position (isotropic) and a certain velocity. There

Method	W_1 (position) ↓	W_1 (rotation) ↓
MC-100	0.0094	1.065
MC-50	0.0147	1.919
MC-10	0.0873	3.537
Surrogate	0.0021	0.792
Local Sampling	0.0021	0.792
Linearization	0.0021	0.781
Sigma-Point	0.0021	1.861

Table 2: Validation of collision-free uncertainty propagation against a Monte Carlo ground truth obtained from 1,000 samples. Reported values are 1-Wasserstein distances measured after 4 s. The first column reports positional error for runs with uncertain initial linear velocity; the second reports rotational error for runs with uncertain initial angular velocity. Lower values indicate better agreement with the ground truth.

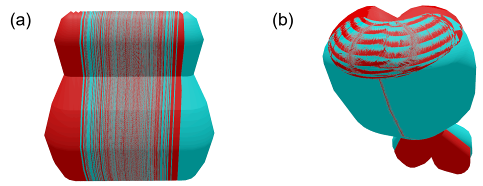

Figure 4: Collision-free comparison between the Monte Carlo ground truth (red, 1,000 samples) and samples drawn from the Gaussian distribution predicted by our linearization method (turquoise, 1,000 samples). (a) Pin with uncertainty in the initial linear velocity. (b) Pin with uncertainty in the initial angular velocity.

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, the method, and positional sigma during the experiments. The models are scaled to fit inside a unit sphere, so the objects are similarly sized, and using the same absolute sigma for all objects is therefore reasonable. We examine scenarios with 10^n spheres, where $n = 1, 2, 3, 4, 5$, corresponding up to 100k spheres. 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 a sampling approach for comparison, where 1k samples serve as the ground truth. A visual comparison after collision can be seen in Fig. 5.

5.2.1 *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. 6 shows the performance. In that plot, we can see that the sampling approach scales linearly with performance for different numbers of samples. On the other hand, it stays almost constant for each sigma and is therefore independent of it. However, our approaches are up to 1-1.5 orders of magnitude faster than ten samples. The surrogate model is the fastest, followed by the local sampling approach and linearization. The linearization approach becomes faster when more spheres are used. The reason for that is that it processes only a single contact point. The sigma-point approach takes the most computation, as expected. Also, the overall computation time increases as the number of spheres increases.

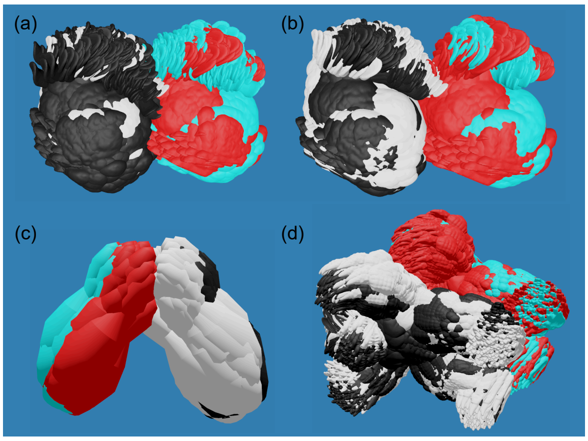

Figure 5: Visual comparison between the ground truth (1,000 Monte Carlo samples; black and red) and our methods for two rigid bodies colliding under initial uncertainty ($\sigma = 0.05$). (a) Bunny with 1,000 spheres using local sampling. (b) Bunny with 1,000 spheres using persistent sigma points. (c) Pin with 100 spheres using linearization. (d) Armadillo with 10,000 spheres using the surrogate model.

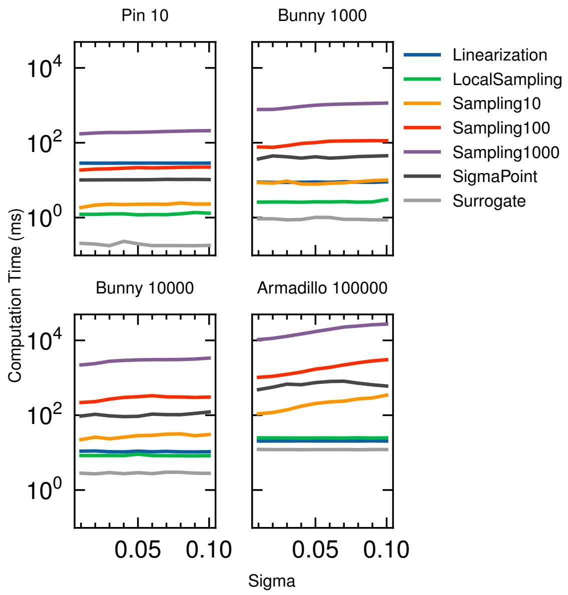

Figure 6: Computation time over sigma for different methods, objects with different numbers of spheres. Our approaches take the least computation time.

5.2.2 *Quality Analysis*

In Fig. 7, we see the average W_1 distance for each sigma, with 1k samples as the ground truth. With more samples, the similarity to the ground truth increases, but not linearly. We can see that across all methods, the W_1 distances increase with larger sigma, as one would expect. Our surrogate model achieves similar accuracy to using 10 samples for position and 5 for rotational accuracy. Linearization and local sampling achieve even better positional accuracy than surrogate, but they are worse at rotational error. The sigma-point approach achieved the best accuracy for position and rotation. We can see that, in general, the relative error with respect to Monte Carlo is larger for rotation than for position. With the shaded regions for 10 and 100 samples (1 standard deviation), we can also see that, for the smallest sigma, the rotational uncertainty is more difficult to represent with Monte Carlo. It is wider than for the position beginning with the smallest sigma. We hypothesize that this is due to the computation of torque, which involves a more significant compounding of uncertainties, thereby making prediction more challenging. On the other hand, this discrepancy might underscore the importance of non-local contact information for accurate angular prediction. However, the sigma-point approach had non-local contacts and achieved rotational accuracy comparable to 5–10 Monte Carlo samples. We think the reason is that we estimate a new mean and covariance from it, whereas the ground-truth rotational distribution is already much less Gaussian-like than for position.

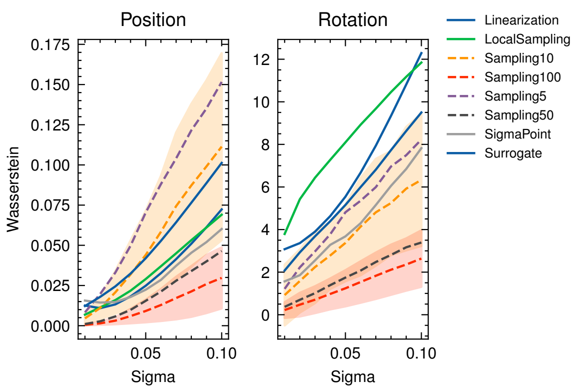

Figure 7: Average W_1 distances to the 1k sampling distribution over sigma for different objects and number of spheres, measured when objects stopped moving after collision. Left: positional W_1 distance (Euclidean norm). Right: rotational W_1 distance (geodesic distance, degrees). Lower values indicate greater similarity to the ground-truth distribution.

Method	Scalar count	Scalars for $B = 2$
Monte Carlo (N samples)	Nn_x	$26N$
Linearization	$n_x + n_\delta^2$	$26 + 24^2 = 602$
Surrogate	$n_x + B \cdot 12^2$	$26 + 2 \cdot 12^2 = 314$
Persistent sigma points	$(2n_\delta + 1)n_x$	$49 \cdot 26 = 1274$
Local sampling	$n_x + B \cdot 12^2$	$26 + 2 \cdot 12^2 = 314$

Table 3: Persistent storage requirements, measured as the number of scalar values required to represent the current simulated distribution. Here B is the number of rigid bodies, $n_x = 13B$ is the dimension of the stored rigid-body state, and $n_\delta = 12B$ is the dimension of the local error state used for covariance propagation.

5.2.3 *Storage Analysis*

We compare persistent storage in terms of the number of scalar values required to represent the current simulated distribution in Table 3. This excludes static geometry data, such as sphere hierarchies, because it is shared by all methods. Let B be the number of rigid bodies, $n_x = 13B$ the dimension of the stored rigid-body state, and $n_\delta = 12B$ the dimension of the local error state used for covariance propagation. The difference arises because orientations are stored as quaternions in the mean state, while rotational uncertainty is represented in the tangent space. The counts assume dense matrix storage.

For the two-body collision scenarios considered in our evaluation, the proposed methods require less persistent storage than high-sample Monte Carlo baselines such as MC-100 and MC-1000. For example, MC-1000 stores 26,000 scalar values, compared with 602 for linearization, 314 for the square-root surrogate and local sampling, and 1274 for persistent sigma points.

More generally, per-body methods such as the square-root surrogate and local sampling scale linearly with the number of rigid bodies and remain more storage-efficient than Monte Carlo for sample counts larger than approximately 12. In contrast, methods that store joint uncertainty, such as linearization with a full covariance and persistent sigma points, scale quadratically with the number of bodies. For sufficiently many bodies and a fixed Monte Carlo sample count, these joint representations can therefore require more storage than Monte Carlo.

6 DISCUSSION

Our results show that uncertainty can be propagated through a rigid-body collision simulation substantially faster than explicit Monte Carlo sampling, while also achieving a competitive accuracy. The comparison highlights a practical trade-off between accuracy and efficiency: persistent sigma points provide the best overall approximation quality, whereas the surrogate model offers the lowest runtime and remains competitive, especially for rotational uncertainty.

Several modeling choices were made to isolate the core problem of uncertainty propagation with rigid-body collisions. We consider penalty-based rigid-body dynamics with linear damping, but omit gravity, floor contacts, and other external effects. The collision experiments focus on pairwise object interactions with controlled initial uncertainty, enabling a systematic comparison of propagation methods but limiting the direct generality of the reported results. Gravity itself is not expected to fundamentally alter the collision-free propagation of uncertainty, but it may increase the frequency and complexity of subsequent contacts. Floor interactions, repeated impacts, stacking, and dense clutter would introduce longer contact sequences and stronger sensitivity to contact timing, and are therefore likely to challenge the proposed approaches more severely than the two-body collisions studied here.

The treatment of uncertainty during collisions also differs across methods. The linearization approach retains only the dominant contact per rigid-body pair to reduce covariance inflation. In contrast, persistent sigma points run through the full simulator and can capture non-local contact effects. This distinction may help explain the sigma-point method's better rotational accuracy. Rotational uncertainty is particularly challenging because torque depends on both force variation and uncertain lever arms, causing approximation errors to compound more strongly than for translation. This interpretation is consistent with the quality analysis, which shows that rotational Wasserstein distances are generally larger than positional ones, and that sigma points perform best overall.

The surrogate model is designed as a low-cost and stable approximation for capturing collision related uncertainty growth. Its update is bounded by the local collision-response scale, as discussed in Sec. 4.2, thereby avoiding excessive uncertainty growth from a single contact event. However, the model approximates the effect of uncertainty on the contact normal direction and does not explicitly account for cross-body correlations. Nevertheless, the surrogate remains competitive with low-sample Monte Carlo baselines while being the fastest method in our evaluation, as shown in Figs. 6 and 7.

More broadly, this work should be viewed as a first step toward a rigorous, uncertainty-aware rigid-body collision simulation. The design space is large, and our experiments suggest that more expressive representations do not automatically yield better approximations. For instance, preliminary variants of local sampling with joint cross-body covariance, although seemingly more principled, were less stable and did not consistently improve agreement with Monte Carlo sampling. This reflects the difficulty of propagating Gaussian uncertainty through collisions.

Scalability and long-term stability are additional limitations. The storage analysis shows that per-body methods, such as the surrogate model and local sampling, scale linearly with the number of rigid bodies. Linearization with a full joint covariance scales quadratically, while the persistent sigma-point method also has quadratic storage in our joint-state formulation, because both the number of sigma points and the size of each propagated state grow linearly with the number of bodies. This suggests that joint representations may become costly in larger scenes, despite their ability to preserve correlations. Over longer horizons, repeated collision events may also amplify model mismatch and accumulate errors in mean and covariance estimates. In general, evaluating stability over longer simulations and testing scenarios with repeated contact transitions would provide a stronger assessment of robustness. This is particularly relevant in collision scenarios, where the resulting distribution may become non-Gaussian after contact. Nevertheless, our experiments show that, for the controlled scenarios considered here, a single Gaussian representation can still provide good agreement with Monte Carlo. Extending the framework to include mixture models or adaptive sampling strategies may further improve accuracy when a single Gaussian representation becomes insufficient.

7 APPLICATION TO ROBOTICS

The proposed uncertainty propagation methods are particularly relevant for autonomous agents operating in unstructured environments. When an embodied agent interacts during everyday activities, its sensory observations are inherently limited temporally (e.g., will the placed object fall?), spatially (e.g., where is the handle?), and informationally (e.g., no depth for transparent object). This sparsity increases the semantic entropy within bottom-up data-driven approaches, often leading to brittle behaviors. Insights from human cognition suggest that humans overcome these limitations by utilizing mental simulation, emulating the hidden dynamics of manipulated objects through an internal physics engine [ZGF+20]. However, as we have

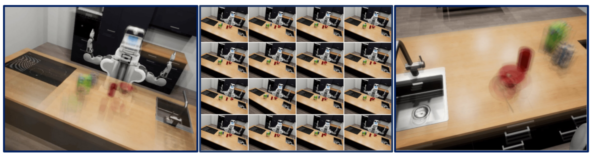

Figure 8: Uncertain simulation enabling safe and successful robot manipulation. Probabilistic anticipation of object states after self-interactions and with the robot.

demonstrated, such mental simulations are only reliable if they account for uncertain physical parameters that cannot be directly observed. By embedding positional and contact uncertainty into these simulations, the agent can leverage the resulting physical knowledge. Our fast uncertainty-aware simulation allows the robot to anticipate the behavior of objects after interactions, be it object-object or object-robot in realtime, and then estimate the pose of these objects, which is essential for safe and successful grasping [KK25].

8 CONCLUSION AND FUTURE WORK

We presented four methods for rigid-body simulation with positional and rotational uncertainty and compared them against Monte Carlo sampling, which is computationally expensive but commonly used as a baseline. Our results show that the proposed methods are significantly faster while still achieving competitive accuracy. Among them, the surrogate model is the fastest and performs well for rotational uncertainty, whereas the sigma-point method gives the best overall accuracy. Linearization and local sampling also achieve strong results, especially for position. Overall, the results suggest that uncertainty can be handled directly during rigid-body collision simulation without relying on a large number of costly samples. In future work, it would be interesting to combine the speed of the surrogate model with the positional accuracy of local sampling, and to explore richer uncertainty representations such as multi-modal distributions or learned models for rotational uncertainty prediction.

9 ACKNOWLEDGMENTS

This research has been (partially) supported by the German Research Foundation DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 "EASE - Everyday Activity Science and Engineering"

REFERENCES

[AMGC02] M Sanjeev Arulampalam, Simon Maskell, Neil Gordon, and Tim Clapp. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. *IEEE Transactions on signal processing*, 50(2):174–188, 2002.

[BHW96] Ronen Barzel, John R Hughes, and Daniel N Wood. Plausible motion simulation for computer graphics animation. In *Computer Animation and Simulation'96: Proceedings of the Eurographics Workshop in Poitiers, France, August 31–September 1, 1996*, pages 183–197. Springer, 1996.

[FCG+21] Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.

[FVCC+24] Rémi Flamary, Cédric Vincent-Cuaz, Nicolas Courty, Alexandre Gramfort, Oleksii Kachaiev, Huy Quang Tran, Laurène David, Clément Bonet, Nathan Cassereau, Théo Gnassounou, Eloi Tanguy, Julie Delon, Antoine Collas, Sonia Mazelet, Laetitia Chapel, Tanguy Kerdoncuff, Xizheng Yu, Matthew Feickert, Paul Krzakala, Tianlin Liu, and Eduardo Fernandes Montesuma. Pot python optimal transport (version 0.9.5), 2024.

[GJ22] Purvi Goel and Doug L James. Unified many-worlds browsing of arbitrary physics-based animations. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022.

[GSS93] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, 140:107–113, 1993.

[HWFS13] Christoph Hertzberg, René Wagner, Udo Frese, and Lutz Schröder. Integrating generic sensor fusion algorithms with sound state representations through encapsulation of manifolds. *Information Fusion*, 14(1):57–77, 2013.

[JU97] Simon J. Julier and Jeffrey K. Uhlmann. New extension of the Kalman filter to nonlinear systems. In Ivan Kadar, editor, *Signal Processing, Sensor Fusion, and Target Recognition VI*, volume 3068, pages 182 – 193. International Society for Optics and Photonics, SPIE, 1997.

[JV00] Johan Jansson and Joris SM Vergeest. A general mechanics model for systems of deformable solids. In *Proceedings of TMCE*, 2000.

[Kal60] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35, 1960.

[KJPZJ24] Nathan J. Kong, J. Joe Payne, James Zhu, and Aaron M. Johnson. Saltation matrices: The essential tool for linearizing hybrid dynamical systems. *Proceedings of the IEEE*, 112(6):585–608, 2024.

[KK25] Franklin Kenghagho Kenfack. *Prospective perception through cognitive emulation for robot manipulation tasks: "Perceiving like humans do"*. Doctoral dissertation, University of Bremen, Bremen, Germany, 2025.

[KNM+22] Franklin K Kenghagho, Michael Neumann, Patrick Mania, Toni Tan, Feroz A Siddiky, René Weller, Gabriel Zachmann, and Michael 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)*, pages 815–822. IEEE, 2022.

[KPCJ21] Nathan J. Kong, J. Joe Payne, George Council, and Aaron M. Johnson. The salted kalman filter: Kalman filtering on hybrid dynamical systems. *Automatica*, 131:109752, 2021.

[LWMC12] Andrew W. Long, Kevin C. Wolfe, Michael Mashner, and Gregory S. Chirikjian. The banana distribution is gaussian: A localization study with exponential coordinates. In *Robotics: Science and Systems*, 2012.

[SHS+24] K. A. Smith, J. B. Hamrick, Adam N. Sanborn, P. W. Battaglia, T. Gerstenberg, T. D. Ullman, and J. B. Tenenbaum. Intuitive physics as probabilistic inference. In T. L. Griffiths, Nick Chater, and Joshua B. Tenenbaum, editors, *Bayesian models of cognition : reverse engineering the mind*. MIT Press, Boston, MA, 2024.

[TJ07] Christopher D Twigg and Doug L James. Many-worlds browsing for control of multibody dynamics. In *ACM SIGGRAPH 2007 papers*, pages 14–es. 2007.

[TMOT12] Min Tang, Dinesh Manocha, Miguel A Otaduy, and Ruofeng Tong. Continuous penalty forces. *ACM Transactions on Graphics (TOG)*, 31(4):1–9, 2012.

[TPBF87] Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. Elastically deformable models. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 205–214, 1987.

[Wit97] Andrew Witkin. Physically based modeling: principles and practice constrained dynamics. *Computer graphics*, 9:27, 1997.

[WZ09] Rene Weller and Gabriel Zachmann. A unified approach for physically-based simulations and haptic rendering. In *Sandbox 2009: ACM SIGGRAPH Video Game Proceedings*, New Orleans, LA, USA, August 2009. ACM Press.

[WZ10] Rene Weller and Gabriel Zachmann. Protosphere: a gpu-assisted prototype guided sphere packing algorithm for arbitrary objects. In *ACM SIGGRAPH ASIA 2010 Sketches*, SA '10, 2010.

[ZGF+20] Yixin Zhu, Tao Gao, Lifeng Fan, Siyuan Huang, Mark Edmonds, Hangxin Liu, Feng Gao, Chi Zhang, Siyuan Qi, Ying Nian Wu, Joshua B. Tenenbaum, and Song-Chun Zhu. Dark, beyond deep: A paradigm shift to cognitive ai with humanlike common sense. *Engineering*, 2020.