Curved Tactile Sensor Simulation with Hydroelastic Contacts in MuJoCo

Florian Patzelt*
CITEC, Faculty of Technology
Bielefeld University
Germany
fpatzelt@techfak.uni-bielefeld.de

David Leins*
CITEC, Faculty of Technology
Bielefeld University
Germany
dleins@techfak.uni-bielefeld.de

Robert Haschke
CITEC, Faculty of Technology
Bielefeld University
Germany
rhaschke@techfak.uni-bielefeld.de

Abstract

This work builds upon prior research on the highly realistic simulation of tactile sensors integrated into MuJoCo, using hydroelastic contact surfaces. Our modifications introduce an additional layer of abstraction, allowing generalization to any sensor shape and overcoming the limitation of exclusively simulating flat surfaces. Using a fingertip sensor as an example, we demonstrate how our extension is able to successfully simulate curved sensor surfaces.

1 Introduction

The sense of touch is essential for both humans and robots, as it allows them to gather information about the world surrounding them. Sensing and interpreting tactile feedback is critical for tasks ranging from delicate object manipulation to safe human-robot collaboration. Robots are becoming more sensitive and dexterous thanks to recent advances in tactile sensing technology.

As tactile sensing capabilities continue to grow, simulation is becoming increasingly important in advancing the field. Just as simulation has helped researchers test and refine algorithms in general robotics research, it is now becoming a cornerstone in the development and application of tactile sensing systems (1; 2; 3). Simulation provides a controlled and cost-effective environment in which researchers can explore, analyze, and fine-tune the nuances of tactile perception without the need for physical prototypes.

In our previous work (4), we have implemented a realistic simulation of tactile sensors in MuJoCo (5) that uses hydroelastic contact surfaces (6). This departure from the conventional point contact model in favor of hydroelastic contact surfaces has enabled them to accurately replicate measurements obtained from a barometric sensor array within the simulation environment.

Hydroelastic contact surfaces constitute an approximate model for calculating the contact area, pressure distribution, and net contact wrench during soft object collisions. They combine the principles of soft-body dynamics and hydroelastic pressure, facilitating the computation of continuous pressure fields within interacting objects. The contact surface is defined at the equilibrium of the overlapping pre-computed pressure fields of the objects. By incorporating dissipative, rate-dependant

*Equal Contribution

NeurIPS 2023 Workshop on Touch Processing: a new Sensing Modality for AI .

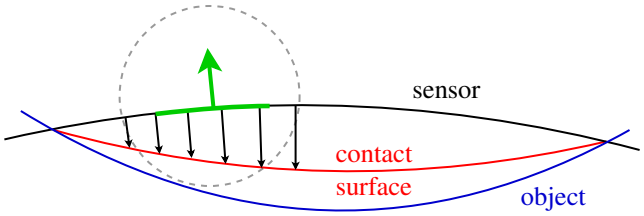

Figure 1: Concept of raycasting, illustrated for the cross-section of a taxel: The black line depicts the surface of the tactile sensor, the blue line the penetrating object, and the red line the resulting contact surface. The green line section and the green arrow depict an individual taxel and its normal located at the center of the taxel's receptive field, which in turn is depicted by the dashed circle. The black arrows represent rays that are cast from the receptive field onto the contact surface.

pressure and friction, continuous wrench values between colliding objects can be calculated from the integral of traction contributions over the contact surface. This approach provides faster computational evaluations than finite element approaches while accurately capturing the force, moment, and stiffness variations associated with soft contacts. In addition, hydroelastic contact surfaces allow pressure to be calculated at any point on the contact surface, simplifying the simulation of tactile sensors.

We have integrated hydroelastic contact surfaces and virtual sensors into our MuJoCo-ROS framework. This framework provides an intuitively accessible interface for plugin integration and facilitates the seamless transition from simulation to real-world applications through its ROS interface, effectively emulating a genuine robotic system.

However, our sensor implementation was limited to a flat, regular taxel array, i.e. requiring a flat, rectangular sensor surface. To address this limitation, we extend our work by introducing a new sensor implementation that employs the same measuring principles but can be effectively applied to arbitrary sensor surfaces. We illustrate the methodology for adapting the approach without additional computational overhead by demonstrating the simulation of a piezo-resistive fingertip sensor for the Shadow Dexterous Hand (7).

2 Methods

In simulation, a sensor comprised of multiple sensor cells (or *taxels*), is represented as a compliant object whose shape is defined by a triangle mesh. The restriction to flat sensors in our previous work allows for a simple definition of the sensor layout. A flat 3D sensor surface effectively boils down to a 2D sensor plane, where the cell normals are given by the normal of the plane and the cell center positions can be inferred by a regular grid with specified width and height. The extension to arbitrary shapes requires us to specifically define the taxel centers on or near to the sensor surface and the normal directions to configure which direction each taxel is most sensitive to. Instead of specifying physical sensor extents, we opt for a receptive field defined by a radius around the cell center to constrain the area in which contacts contribute to the output of the sensor cell. This reflects the fact that forces are distributed isometrically on the sensor, regardless of the exact shape and size of the taxel.

When the sensor comes into contact with an object, we first compute the hydroelastic contact surface, which is represented as a triangle mesh with associated pressure values in each vertex. Using barycentric coordinates, we can interpolate the pressure at any point on the surface. For more details on the computation of hydroelastic contact surfaces and pressure values, we refer to (6) and (8).

Similar to our previous work, we determine tactile sensor readings by mapping pressure values from the hydroelastic contact surface onto the individual taxels of the sensor. The sensor reading of a taxel resembles the integral of distance-weighted pressure values within the taxel's receptive field. We approximate this integral by sampling discrete points on the sensor surface within the taxel's receptive field, projecting them onto the contact surface, and aggregating the pressure values at these points.

The projection onto the contact surface is framed as a rendering problem: taxels are treated as pixels for which rays are cast opposite to the taxel's surface normal. The corresponding points on

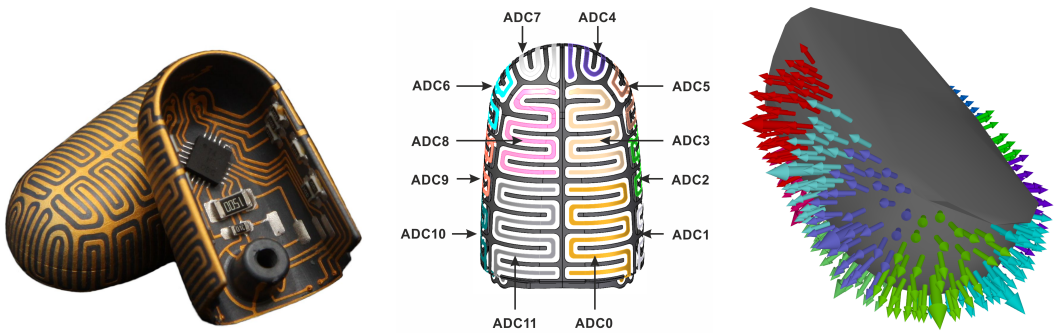

Figure 2: Fingertip of the Shadow Robot Hand (7) (left), its taxel layout (center), and the simulated sensor (right). Large arrows indicate the center and normal of individual taxels. Small arrows indicate the surface normal at the sampling points. The color indicates the nearest taxel to which they are assigned.

the contact surface can then be determined as the first intersection of these rays with the contact surface. Following our previous work, we leverage techniques from the field of computer graphics to efficiently solve the ray-triangle intersection problem. We reduce the number of intersection tests by employing bounding volume hierarchies (BVH), specifically Wald's binned building approach (9). For each contact surface on a sensor, we construct a bottom-level acceleration structure (BLAS) based on axis-aligned bounding boxes and combine them in top-level acceleration structures (TLAS). These strategies significantly reduce the computational overhead associated with intersection tests, enabling the simulation of reasonably sized sensors in real-time. A 2D scheme of the raycasting procedure is depicted in Figure 1.

Various aggregation functions can be considered to compute the pressure output of each taxel from the collected pressure samples. We found that the convex weighting scheme we employed in our previous work, which decreases quadratically with the distance from the taxel center, works well for curved sensors:

$$p = \frac{\sum_{i=0}^{n-1} w_i p_i}{\sum_{i=0}^{n-1} w_i} \qquad w_i = (r - d_i)^2$$

Here, n denotes the number of rays cast, p_i represents the pressure at the point where ray i intersects the contact surface, r signifies the radius of the taxel's receptive field, and d_i is the distance between the center of the taxel and the origin of ray i on the sensor surface.

The main challenge in generalizing from flat sensor surfaces to curved sensor surfaces lies in finding suitable sample points on the sensor surface within the receptive field of each taxel. To get a good approximation of the pressure integral over the corresponding patch of the contact surface, the samples must be uniformly distributed over the sensor surface and not clustered in one spot. For flat sensor arrays, this is trivial to achieve. However, this is a considerable issue for arbitrarily shaped sensor surfaces. To overcome this, we use Constrained Poisson-disk sampling (10) on the sensor surface mesh as a preprocessing step. This gives us a set of points that are uniformly distributed over the sensor surface, i.e. the distance between each point and its nearest neighbor is maximal. A resolution parameter defines how many samples should be generated per surface area of the sensor. Similar to the resolution parameter employed in our previous work, this parameter allows us to balance between computational load and accuracy of the sensor output.

We assign sample points to taxels based on whether they lie within a taxel's receptive field. Furthermore, we found it sensible to constrain the angle between the taxel's normal and the surface normal at a sampling point to be less than $45°$. This emulates the effect that real tactile sensors usually have a specific origin and normal for which they are most sensitive. Deviating from either reduces the sensor's response. Note that due to overlapping receptive fields, samples may be assigned to multiple adjacent taxels. This mimics the behavior of a real sensor in terms that pressure at a certain point on the sensor surface might cause a sensor response in multiple adjacent taxels. The sample points determined by this method for a virtual fingertip of the Shadow Dexterous Hand are visualized in

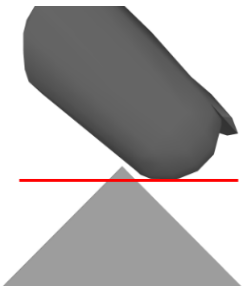

Figure 3: We use a soft constraint to move the fingertip parallel to the red line, sliding it over the edge.

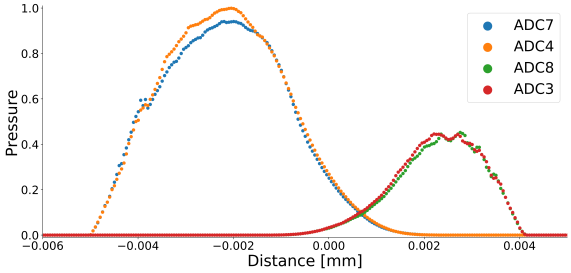

Figure 4: The sensor response for the four taxels that came into contact with the object's edge. The pressure was normalized using the maximum recorded pressure.

Figure 2. All sample positions and taxel assignments can be computed before the simulation starts, and only samples that are assigned to at least one taxel are used. Thus, these computations do not cause any additional simulation overhead.

3 Experiment

To validate our sensor implementation, we simulate the 12-taxel fingertip sensor for the Shadow Robot Hand (7) moving across the edge of an object as depicted in Figure 3. A soft constraint is used to guide the fingertip and press it against the edge. The sensor output is recorded every 0.1 mm. Only four taxels come into contact with the edge. The sensor values of these taxels are visualized in Figure 4.

Similar to a real sensor, the sensor values increase and decrease smoothly as the fingertip slides over the edge. Note that the taxels at the tip of the finger have much more contact with the edge than the taxels further back, resulting in higher pressure values.

While examining the sensor values, we noticed some artifacts that look similar to we described in our previous work (4). We found that at some stages of the experiment where the pressure integral over the entire contact surface remains largely the same, the pressure integral over patches of the contact surface changes rapidly. For example, a contact surface with the same pressure integral can have the pressure evenly distributed over all triangles, or some small triangles with high pressure and some larger triangles with less pressure. While all of these distributions result in the same contact force, this in combination with the sampling-based nature of our approach causes the artifacts.

4 Conclusion and future work

We have extended the simulated sensor implementation proposed in our previous work (4) to be applicable to arbitrarily shaped sensor surfaces. For a virtual fingertip sensor, we demonstrate that raycasting positions are well distributed over the curved sensor surface and that sensor readings smoothly increase and decrease as the fingertip sensor is moved over an edge.

In the future, we plan to conduct a qualitative study comparing our simulated fingertip to the real fingertip in order to validate the plausibility of the measured sensor values. Further, we consider increasing the accuracy of our sensor simulation through multiple measures:

In order to make the parameterization more independent from the curvature of the sensor surface, the geodesic distance can be used instead of the Euclidean distance to determine which samples on the sensor surface are in the receptive field of a taxel.

To better emulate the real structure of the sensor cells, taxels can be decomposed into multiple subtaxels defined by separate centroids, normals, and (smaller) receptive fields. For a taxel, the set of all samples on the sensor surface over all subtaxels is used. During aggregation, the weighting can be computed considering the distance to the nearest subtaxel.

Acknowledgments and Disclosure of Funding

This work was supported by the BMBF project Sim4Dexterity.

References

[1] Z. Si and W. Yuan, "Taxim: An example-based simulation model for gelsight tactile sensors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2361–2368, 2022.

[2] S. Wang, M. Lambeta, P.-W. Chou, and R. Calandra, "Tacto: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3930–3937, 2022.

[3] Y. Wang, W. Huang, B. Fang, F. Sun, and C. Li, "Elastic tactile simulation towards tactile-visual perception," in *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, (New York, NY, USA), p. 2690–2698, Association for Computing Machinery, 2021.

[4] D. Leins, F. Patzelt, and R. Haschke, "More accurate tactile sensor simulation with hydroelastic contacts in mujoco," in *IROS 2023 Workshop on Leveraging Models for Contact-Rich Manipulation*, 2023.

[5] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033, IEEE, 2012.

[6] R. Elandt, E. Drumwright, M. Sherman, and A. Ruina, "A pressure field model for fast, robust approximation of net contact force and moment between nominally rigid objects," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 8238–8245, IEEE, 2019.

[7] R. Kõiva, M. Zenker, C. Schürmann, R. Haschke, and H. J. Ritter, "A highly sensitive 3d-shaped tactile sensor," *IEEE/ASME Int. Conf. on Advanced Intelligent Mechatronics*, pp. 1084–1089, 2013.

[8] J. Masterjohn, D. Guoy, J. Shepherd, and A. Castro, "Velocity level approximation of pressure field contact patches," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11593–11600, 2022.

[9] I. Wald, "On fast construction of sah-based bounding volume hierarchies," in *2007 IEEE Symposium on Interactive Ray Tracing*, pp. 33–40, IEEE, 2007.

[10] M. Corsini, P. Cignoni, and R. Scopigno, "Efficient and flexible sampling with blue noise properties of triangular meshes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 6, 2012.