

Hydrosoft: Non-Holonomic Hydroelastic Models for Compliant Tactile Manipulation

Miquel Oller An Dang Nima Fazeli
 Department of Robotics, University of Michigan
 Ann Arbor, MI 48109, United States
 {oller, andang, nfz}@umich.edu
<https://www.mmintl.com/hydrosoft>




Fig. 1. **Approach Overview:** *First panel:* The robot grasps an object using compliant tactile sensors and performs in-hand rotation. *Second panel:* Our method models the tactile sensors as non-holonomic hydroelastic elements. *Third panel:* The non-holonomic hydroelastic model is integrated with differentiable dynamics. *Fourth panel:* The resulting dynamics are used with gradient-based trajectory optimization to control and execute the task.

Abstract—Tactile sensors have long been valued for their perceptual capabilities, offering rich insights into the otherwise hidden interface between the robot and grasped objects. Yet their inherent compliance—a key driver of force-rich interactions—remains underexplored. The central challenge is to capture the complex, nonlinear dynamics introduced by these passive compliant elements. Here, we present a computationally efficient non-holonomic hydroelastic model that accurately models path-dependent contact force distributions and dynamic surface area variations. Our insight is to extend the object’s state space, explicitly incorporating the distributed forces generated by the compliant sensor. Our differentiable formulation not only accounts for path-dependent behavior but also enables gradient-based trajectory optimization, seamlessly integrating with high-resolution tactile feedback. We demonstrate the effectiveness of our approach across a range of simulated and real-world experiments and demonstrate the importance of modeling the path dependence of sensor dynamics.

I. INTRODUCTION

Contact-rich object manipulation is a cornerstone of dexterous robotic systems, enabling precise control over an object’s position and orientation by regulating the forces at the robot–object interface [1], [2]. A particularly effective way to achieve this regulation is through mechanical compliance at the point of contact [3], [4], [5], [6], which can absorb uncertainty and maintain stable force exchange in complex manipulation tasks.

However, capturing the nonlinear, high-dimensional dynamics of compliant end-effectors remains an open problem. On one end of the spectrum, finite element methods can

accurately model deformation but are often computationally prohibitive for real-time control [7]. On the other, simplified spring-damper approximations offer speed but fail to generalize to tasks where the distribution of contact forces and surface deformations play a pivotal role. To bridge this gap, hydroelastic modeling has emerged as an appealing middle ground, balancing physical fidelity with computational efficiency [8], [9]. By focusing on force distributions and contact-surface variations, hydroelastic models bypass much of the complexity associated with fully simulating elastic deformations—making them especially well suited to contact-rich manipulation.

In this work, we propose Hydrosoft: a non-holonomic extension to the hydroelastic model that models: (i) the path dependence of force build up in the sensor substrate; (ii) Coulomb friction; (iii) non-smooth contact transitions. Our approach incorporates relaxations for smooth gradients, offers efficient and parallelized computation, and supports high-resolution tactile sensing for real-time feedback and control. By incrementally updating the contact forces, our method models the history-dependent deformation and force build up of the sensor substrate while preserving the computational benefits of hydroelastic models. As a result, our approach enables gradient-based planning and control in scenarios that demand both precise force regulation and continuous contact ranging from rolling objects on a table to in-hand object reconfiguration.

In more detail, our model addresses 4 key limitations

of prior models: First, the original formulation [8] does not model static friction forces. This is of critical importance when considering grasping and object transitions in the grasp. Second, the hydroelastic formulation does not explicitly model deformation. As a result, sticking shear deformations may result in breaking contact locally between the hydroelastic and rigid objects. Intuitively, when a deformable is in contact and shear is applied, the interface can stick and due to compliance, the contact is preserved. However, this important effect is currently not modeled. Third, the contact forces are solely a function of the relative position between the objects, without accounting for the prior trajectory. Consequently, the computation of contact forces assumes path independence. However, this assumption breaks down in cases involving shear deformations, where the deformation history significantly influences the resulting forces. Finally, the hybrid nature of contact results in the model exhibiting non-smooth dynamics. As a result, the model is only able to produce sub-gradients. Specifically, only the differential elements within the hydroelastic contact region contribute to the gradients, while elements outside the contact region have zero gradients. This limitation is particularly evident in scenarios such as in-hand object rolling, where smooth transitions in contact are critical for accurate force and motion control.

II. RELATED WORK

A. Rigid Body Dynamics with Frictional Contacts

Modeling contact dynamics is challenging due to its hybrid nature, which leads to formulations with complementarity constraints to enforce non-penetration and Coulomb friction. Such problems are often posed as Nonlinear Complementarity Problems (NCPs), but can be approximated by Mixed Linear Complementarity Problems (MLCPs) through polyhedral friction cone approximations [10], [11]. Anitescu et al. [12] reformulated the problem as a Second-Order Cone Program (SOCP), interpretable as KKT conditions of a convex program, enabling gradient propagation [13]. Alternative formulations avoid complementarity: [14] proposed an unconstrained optimization using intermediate variables, while [15] introduced a differentiable closed-form model with spring forces in the contact dual cone. However, these works mainly consider rigid point contacts, whereas our focus is on elastic elements where deformation alleviates interpenetration issues.

B. Smoothing and Differentiability

Contact dynamics are also non-smooth, restricting optimization to sub-gradients. Smoothing methods address this: randomized smoothing [16], log-barrier smoothing [17], [?], and `softplus`-based relaxations [15]. These approaches enable differentiability but may introduce artifacts such as forces-at-a-distance. We similarly apply smoothing to support gradient-based trajectory optimization, while incorporating contact feedback to mitigate model mismatch.

Symbol	Space	Description
m	\mathbb{R}_+	Object Mass
μ	\mathbb{R}_+	Contact Friction
\mathbf{w}_{ext}	\mathbb{R}^6	External Wrench
\mathbf{q}_l^O	$\text{SE}(3)$	Rigid Object Pose
\mathbf{q}_d^H	$\text{SE}(3)$	Hydroelastic Object Pose
\mathbf{q}_d^O	$\text{SE}(3)$	Desired Rigid Object Pose
\mathbf{f}^c	\mathcal{K}_μ^3	Contact Force
\mathbf{F}^c	$\mathbb{R}^{N \times M \times 3}$	Contact Forces
\mathbf{J}^{c_i}	$\mathbb{R}^{6 \times 3}$	Contact Jacobian
\mathbf{M}	$\mathbb{R}^{6 \times 6}$	Generalized Inertia Matrix

TABLE I
NOTATION

C. Modeling Patch Contact

Patch contact has been studied via sliding models [18], [19], extended to precise manipulation [20], [21], though often limited to circular patches and uniform pressure assumptions. Composite friction cones approximate polygonal patches with edge-based cones [22], naturally capturing transitions from patch to line or point contacts. Patch contact has also been modeled for deformable elements, with approaches that estimate contact regions from tactile deformation [23] or jointly recover both deformations and forces [24], [7]. Despite the physical realism, they are often computationally demanding, limiting their use in real-time manipulation.

Hydroelastic models [8], [9] provide an alternative, where patch size emerges from geometry intersections. The Pressure Field Contact (PFC) model [8] is efficient but lacks static forces; [9] extended it to include static friction via optimization, later combined with iLQR for manipulation and locomotion [25], though at high computational cost. We build on hydroelastic formulations for fast, efficient modeling of highly elastic elements to regularize contact size in manipulation.

Contact patch modeling is of special interest in tactile simulation. Differentiable frameworks [26] and libraries like TacSL [27] enable realistic visuotactile outputs, but they often fail to capture tangential forces under sticking contact. In contrast, our formulation produces tangential forces for static friction consistent with Coulomb’s model, improving realism in contact-rich interactions.

III. PROBLEM STATEMENT

We consider the problem of manipulating a rigid object with known geometry and physical parameters from an initial configuration \mathbf{q}_0^O to a desired goal configuration \mathbf{q}_d^O . The robot is equipped with N passive-compliant elements, instantiated here as tactile sensors, which are rigidly attached to it. The robot has full control over the $\text{SE}(3)$ poses of the attachment points, and thus the motion of the compliant elements is implicitly determined by the robot’s configuration. The objective is to compute a sequence of attachment-point motions $\mathbf{q}_1^H, \dots, \mathbf{q}_K^H$ that drive the object toward the goal configuration \mathbf{q}_d^O while respecting the physical constraints arising from the interaction between the robot, the compliant elements, the object, and the environment.




Fig. 2. **Nonholonomic Hydroelastic Virtual Stick-Slip Behaviour:** We show the transition of a differential of the area along the hydroelastic body \mathcal{H} . Starting at a contact-free position \mathbf{p}_0 , i.e. $\phi(\mathbf{p}_0) > 0$, transition 0-1 shows a contact transition where the contact force is proportional to the displacement within the hydroelastic body \mathbf{d}_0 . Note that this creates a virtual reference point $\hat{\mathbf{p}}$ at the boundary of the hydroelastic body. $\hat{\mathbf{p}}$ can be interpreted as the point corresponding to zero contact force in a slip-free transition. Transition 1-2 show a ‘sticking’ transition within the hydroelastic element. Note that contact force stays within the friction cone and the virtual reference point remains the same. Finally, transition 2-3 corresponds to a ‘slipping’ transition, where Coulumb’s friction constraints make the boundary reference point $\hat{\mathbf{p}}$ has virtually translated along the hydroelastic boundary to $\hat{\mathbf{p}}'$

IV. METHODOLOGY

Our methodology combines hydroelastic modeling of compliant contact with trajectory optimization to enable dexterous manipulation. The section is organized as follows.

First, we introduce the hydroelastic formulation used to resolve compliant contact forces between the object and the robot’s tactile elements (Sec. IV-A). We then extend this formulation to a non-holonomic incremental model that captures path-dependent deformations and frictional effects (Sec. IV-B). To improve computational efficiency and differentiability, we further propose a smoothed variant of the hydroelastic model (Sec. IV-C). Next, we embed the contact model into a quasi-dynamic system representation that relates robot control actions—changes in the attachment poses of the compliant elements—to object motion (Sec. IV-D). Finally, we formulate the trajectory optimization problem as a model predictive control (MPC), which computes the sequence of robot actions that achieves the desired object configuration (Sec. IV-E).

A. Compliant Contact Forces via Hydroelastic Models

We will start by introducing the hydroelastic formulation proposed in [8] and [9] adapted for our problem formulated in section III. We will consider the case of resolving the contact forces for an interaction between a rigid object \mathcal{O} and a hydroelastic body \mathcal{H} . In section IV-B we will extend the formulation to the non-holonomic case.

The rigid object \mathcal{O} is discretized into a set of surface points \mathcal{O} where

$$\mathcal{O} = \{\mathbf{p}_1, \dots, \mathbf{p}_N \mid \mathbf{p}_i \in \mathbb{R}^3\} \quad (1)$$

Each point \mathbf{p}_i has associated a surface area A_i and a contact frame centered at the point. We denote as ${}^{\mathcal{O}}\dot{\mathbf{q}}^{c_i}$ as the

configuration of the contact frame $i \in \{1, \dots, N\}$ w.r.t the object frame.

We denote ϕ as the signed-distance function (SDF) of the hydroelastic body for its nominal shape, i.e. contact-free shape. Therefore, we can define the hydroelastic domain as all the points inside the hydroelastic element:

$$\mathcal{H} \triangleq \{\mathbf{p} \in \mathbb{R}^3 \mid \phi(\mathbf{p}) \leq 0\} \quad (2)$$

We seek to compute the contact force associated to each surface point $\mathbf{p}_i \in \mathcal{O}$, denoted as \mathbf{f}^{c_i} . The normal contact force is proportional to the area A_i associated to that point and the penetration distance on the hydroelastic body. The penetration distance is computed using the SDF function. This takes the following expression:

$$f_{n,i} = \max(-EA_i\phi(\mathbf{p}_i), 0) \quad (3)$$

where $E > 0$ is the effective pressure gradient (in [Pa/m]).

To compute the tangential contact forces $\mathbf{f}_t^{c_i}$, [8] assumes them to be 0 if there is no relative motion between the contact point \mathbf{p}_i and the hydroelastic body \mathcal{H} in the tangent plane, i.e. ${}^{\mathcal{O}}\dot{\mathbf{q}}_{\perp}^{c_i} = \mathbf{0}$. Otherwise, it follows Coulumb’s friction.

$$\mathbf{f}_t^{c_i} = \begin{cases} \mathbf{0} \in \mathbb{R}^3 & \text{if } \|{}^{\mathcal{O}}\dot{\mathbf{q}}_{\perp}^{c_i}\| = 0 \\ -\mu f_{n,i} \frac{{}^{\mathcal{O}}\dot{\mathbf{q}}_{\perp}^{c_i}}{\|{}^{\mathcal{O}}\dot{\mathbf{q}}_{\perp}^{c_i}\|} & \text{if } \|{}^{\mathcal{O}}\dot{\mathbf{q}}_{\perp}^{c_i}\| > 0 \end{cases} \quad (4)$$

Similarly, [9] computes the friction force using complementary conditions derived from the maximum dissipation principle. Note that this constraint is combined with other contact constraints and dynamic constraints to resolve the contact via optimization.

Finally, the contact force is computed by combining normal and tangential forces

$$\mathbf{f}^{c_i} = f_{n,i}\hat{\mathbf{n}} + \mathbf{f}_t^{c_i} \quad (5)$$

B. Non-holonomic Incremental Hydroelastic Model

Our goal is to derive a model capable of computing the reaction forces resulting from the contact interaction between a hydroelastic body and a rigid object. Our main considerations are the deformation path, surface area variations, and obeying Coulumb’s friction. To this end, we extend the hydroelastic formulation from [8]. Here, we derive the equations for a single hydroelastic body, but they can trivially be extended to consider M bodies.

We start off with the standard Equations of Motion (EoM) for our system:

$$\mathbf{M}_H(\mathbf{q}^H)\ddot{\mathbf{q}}^H + \mathbf{C}_H(\mathbf{q}^H, \dot{\mathbf{q}}^H)\dot{\mathbf{q}}^H = \tau_g^H + \mathbf{w}_{\text{ext}}^H + \mathbf{u} \quad (6)$$

$$\mathbf{M}_O(\mathbf{q}^O)\ddot{\mathbf{q}}^O + \mathbf{C}_O(\mathbf{q}^O, \dot{\mathbf{q}}^O)\dot{\mathbf{q}}^O = \tau_g^O + \mathbf{w}_{\text{ext}}^O \quad (7)$$

where $\mathbf{w}_{\text{ext}}^H$ and $\mathbf{w}_{\text{ext}}^O$ are external wrenches computed from our hydroelastic contact model, \mathbf{M} denotes the inertia matrix, \mathbf{C} denotes the Coriolis and Centrifugal forces, τ_g the torques induced by gravity, and \mathbf{u} denotes the control input.

In our hydroelastic formulation, we individually compute the contact forces \mathbf{f}^{c_i} for each contact point on the rigid object \mathbf{p}_i . Instead of calculating reaction forces through

inter-object penetration from the closest boundary point, we calculate the reaction forces by tracking the in-penetration displacements. This formulation approximates membrane deformation over time and provides a more accurate model for modelling the deformable object dynamics without accruing significant computational cost.

We denote E and K as the elastic modulus in the normal direction and tangential direction respectively. With this, we can calculate the accrued friction force with the equation:

$$\dot{\mathbf{f}}^{c_i} = H(-\phi(\mathbf{q}^H, \mathbf{q}^O)) \begin{bmatrix} EA_i & 0 & 0 \\ 0 & KA_i & 0 \\ 0 & 0 & KA_i \end{bmatrix} \mathbf{R}_{c_i}(\mathbf{q}^O)^\top \dot{\mathbf{q}}^O \quad (8)$$

where H is the Heaviside step function and $\mathbf{R}_{c_i}^\top$ rotates $\dot{\mathbf{q}}^O$ to be in the contact frame. We add H to ensure we only accumulate our reaction forces when the object is inside \mathcal{H} .

Similar to [14] we enforce the Coulomb Friction Law by projecting the contact force \mathbf{f}^{c_i} onto the friction cone:

$$\tilde{\mathbf{f}}^{c_i} = \begin{bmatrix} \text{ReLU}(\mathbf{f}_n^{c_i}) \\ \min(1, \frac{\mu \max(\mathbf{f}_n^{c_i}, 0)}{\|\mathbf{f}_t^{c_i}\|}) \mathbf{f}_t^{c_i} \end{bmatrix} \quad (9)$$

We also model a force reset which ensures the forces are only active when the object resides in \mathcal{H} , allowing forces to be reset when contacts transition out from \mathcal{H} .

$$\tilde{\mathbf{f}}^{c_i} = H(-\phi(\mathbf{q}^H, \mathbf{q}^O)) \tilde{\mathbf{f}}^{c_i} \quad (10)$$

We finally sum the wrench applied by each $\tilde{\mathbf{f}}^c$ and apply them to both the hydroelastic and rigid objects.

$$\mathbf{w}_{\text{ext}}^H = \sum_{i=1}^{N \times M} \mathbf{H} \mathbf{J}^{c_i \top} \tilde{\mathbf{f}}^{c_i} \quad (11)$$

$$\mathbf{w}_{\text{ext}}^O = \sum_{i=1}^{N \times M} \mathbf{O} \mathbf{J}^{c_i \top} \tilde{\mathbf{f}}^{c_i} \quad (12)$$

$$\mathbf{F}^c = [\tilde{\mathbf{f}}^{c_1} \quad \tilde{\mathbf{f}}^{c_2} \quad \dots \quad \tilde{\mathbf{f}}^{c_{N \times M}}]^\top \quad (13)$$

We wish to discretize our dynamics and will specifically focus on discretizing $\tilde{\mathbf{f}}^c$. The most straightforward discretization using first-order approximation gives the following:

$$\mathbf{f}_{k+1}^{c_i} = \mathbf{f}_k^{c_i} + \dot{\mathbf{f}}_k^{c_i} h$$

where h is the time-step. However, this discretization is not informative for in-contact and out-of-contact transition since it assumes that the rigid object will always be inside or outside of \mathcal{H} from k to $k+1$. This can introduce errors that can propagate and increase over time. To increase the accuracy of our model, we seek a more accurate discretization of in-penetration and out-of-penetration transitions. In particular, we will compute the in-penetration fraction of $\mathbf{f}^{c_i, k}$ from k and $k+1$ using ϕ .

First, for each differential of area, we compute their displacement in \mathcal{H} , denoted by \mathbf{d} . This corresponds to the deformation of the hydroelastic element during that step.




Fig. 3. **Dynamics Computation Graph:** Showing the dependencies to compute next state \mathbf{x}_{k+1} from previous state \mathbf{x}_k and control \mathbf{u}_k .

$$\alpha_d = \frac{\text{ReLU}(-\phi_{k+1}) - \text{ReLU}(-\phi_k)}{\phi_k - \phi_{k+1}} \quad (14)$$

$$\mathbf{d} = \alpha_d \mathbf{r} = \alpha_d (\mathbf{p}_k - \mathbf{p}_{k+1}) \quad (15)$$

where $\phi_k = \phi(\mathbf{p}_k)$. Note that α_d corresponds to the fraction of the hydroelastic element displacement \mathbf{r} inside the hydroelastic domain \mathcal{H} for the . If both \mathbf{p}_k and \mathbf{p}_{k+1} are outside the hydroelastic element, i.e. $\phi_k, \phi_{k+1} \geq 0$, then $\alpha_d = 0$. Furthermore, if they are both in \mathcal{H} , i.e. $\phi_k, \phi_{k+1} \leq 0$, then $\alpha_d = 1$.

Next, we project \mathbf{d} into normal and tangential components of the contact frame. $d_n = \langle \mathbf{d}, \hat{\mathbf{n}} \rangle$ is the projection of \mathbf{d} along the normal direction of the contact frame $\hat{\mathbf{n}}$. $\mathbf{d}_t = \mathbf{d} - d_n \hat{\mathbf{n}}$ is the projection of \mathbf{d} in the contact plane. We update the normal forces using the normal displacement d_n and the normal elastic modulus E , and tangential forces with tangential displacement \mathbf{d}_t and tangential elastic modulus K .

$$f_{n, k+1}^{c_i} = f_{n, k}^{c_i} + EA_i d_n \quad (16)$$

$$\mathbf{f}_{t, k+1}^{c_i} = \mathbf{f}_{t, k}^{c_i} + KA_i \mathbf{d}_t \quad (17)$$

We then enforce the friction cone constraint.

$$\tilde{f}_{n, k+1}^{c_i} = \text{ReLU}(f_{n, k+1}^{c_i}) \quad (18)$$

$$\tilde{\mathbf{f}}_{t, k+1}^{c_i} = \min(1, \frac{\mu \tilde{f}_{n, k+1}^{c_i}}{\|\mathbf{f}_{t, k+1}^{c_i}\|_2}) \mathbf{f}_{t, k+1}^{c_i} \quad (19)$$

Finally, we calculate force reset if the rigid object exits \mathcal{H} :

$$\tilde{\mathbf{f}}_{k+1}^{c_i} = H(-\phi_{k+1}) (\tilde{f}_{n, k+1}^{c_i} \hat{\mathbf{n}} + \tilde{\mathbf{f}}_{t, k+1}^{c_i}) \quad (20)$$

C. Smoothed Hydroelastic Model

To enhance the model's computational efficiency and differentiability, we introduce smoothing by relaxing the non-holonomic incremental hydroelastic model. Specifically, we set $\alpha_d = 1$ and replace the Heaviside step function H with a sigmoid_{β} function. This modification weights the force contributions based on the displacement's position relative to the hydroelastic domain \mathcal{H} . Displacements outside \mathcal{H} contribute minimally, while those inside the domain contribute fully, ensuring a smooth transition at the boundary. This approach mitigates abrupt changes in force computations making the model suitable for gradient-based optimization.

D. Quasi-dynamic Formulation

Given the path-dependent hydroelastic model described in section IV-B, we can integrate it to recover the dynamics. Given the control actions $\mathbf{u} \in \text{SE}(3)$ as changes of poses of the hydroelastic bodies positions, the dynamics take the following form

$$\mathbf{x}_{k+1} = f_{\text{dyn}}(\mathbf{x}_k, \mathbf{u}_k)$$

where the state is composed by the object pose, hydroelastic bodies pose, and contact forces:

$$\mathbf{x}_k = [\mathbf{q}_k^{\text{O}} \quad \mathbf{q}_k^{\text{H}} \quad \mathbf{F}^{\text{c}_k}]^{\top}$$

We adopt a quasi-dynamic integration scheme, following prior work where it has been shown to be effective for modeling contact-rich interactions [28], [17]. The procedure consists of the following steps, illustrated in Fig. 3:

- 1) Apply control actions.

$$\mathbf{q}_{k+1}^{\text{H}} = \mathbf{q}_k^{\text{H}} + \mathbf{u}_k$$

- 2) Aggregate contact forces

$$\mathbf{w}_{\text{ext}}^{\text{O}} = \sum_{i=1}^{N \times M} \mathbf{O} \mathbf{J}_i^{\text{c}_i \top} \tilde{\mathbf{f}}^{\text{c}_i}$$

- 3) Integrate object dynamics

$$\mathbf{q}_{k+1}^{\text{O}} = \mathbf{q}_k^{\text{O}} + \frac{h^2}{\epsilon} [\mathbf{M}^{-1}(\mathbf{w}_{\text{ext}}^{\text{O}} + \boldsymbol{\tau}_g)]$$

- 4) Compute next contact forces

$$\mathbf{F}^{\text{c}_{k+1}} = \text{compute_forces}(\mathbf{q}_{k+1}^{\text{O}}, \mathbf{q}_{k+1}^{\text{H}}, \mathbf{q}_k^{\text{O}}, \mathbf{q}_k^{\text{H}}, \mathbf{F}^{\text{c}_k})$$

E. Control Formulation

In this section, we present the trajectory optimization formulation, which integrates the proposed hydroelastic model with model predictive control (MPC). The objective is to determine a sequence of control actions $\mathbf{u} \in \text{SE}(3)$ that minimizes a prescribed cost function while satisfying the system dynamics governed by f_{dyn} .

$$\begin{aligned} \min_{\mathbf{u}_{0:K-1}} \quad & J = \Phi(\mathbf{x}_K) + \sum_{k=0}^{K-1} \ell_k(\mathbf{x}_k, \mathbf{u}_k) \\ \text{s.t.} \quad & \mathbf{x}_{k+1} = f_{\text{dyn}}(\mathbf{x}_k, \mathbf{u}_k), \quad k = 0, \dots, K-1, \\ & \mathbf{x}_k \in \mathcal{X}_k, \quad \mathbf{u}_k \in \mathcal{U}_k, \quad k = 0, \dots, K-1, \\ & \mathbf{x}_0 = \mathbf{x}_{\text{init}}. \end{aligned} \quad (21)$$

Here, $\ell_k(\cdot)$ denotes the stage cost, and $\Phi(\cdot)$ the terminal cost, both assumed to be differentiable.

Since the dynamics and cost are differentiable, we employ the gradient-based trajectory optimization method from [29]. This approach integrates Cross Entropy Methods (CEM) with gradient descent, mitigating issues with poor initialization and the locality of purely gradient-based optimization. The procedure is summarized in Algorithm 1.

To explore a broad solution space, we parallelize over B trajectories at each iteration. The best M trajectories

(lowest cost) are used to update the control distribution parameters, while the lowest-cost trajectory determines the control applied to the system.

At each time step, sensor feedback and state estimation provide an updated initial condition, from which the optimization is re-solved. This receding-horizon strategy introduces feedback, enabling adaptation to disturbances, modeling errors, and environmental uncertainty.

Algorithm 1: Iterative Cross-Entropy Gradient-Based Trajectory Optimization

Input : Reference distribution parameters

$\mu_{0:K-1}, \Sigma_{0:K-1}$, dynamics f , cost function J , number of parallel trajectories B , gradient descent iterations G , elite set size $M \leq B$.

Output: Optimized control sequence $\mathbf{u}_{0:K-1}^*$.

```

1 for each time step  $h = 0$  to  $H - 1$  do
2   Measure current state  $\mathbf{x}_h$ ;
3   Sample  $B$  control sequences  $\{\mathbf{u}_{h:K-1}^{(b)}\}_{b=1}^B$  from
    $\mathcal{N}(\mu_{h:K-1}, \Sigma_{h:K-1})$ ;
4   for each gradient iteration  $i = 1$  to  $G$  do
5     for each trajectory  $b = 1$  to  $B$  do
6       Roll out trajectory  $\{\mathbf{x}_k^{(b)}\}_{k=h}^K$  with
       dynamics  $\mathbf{x}_{k+1}^{(b)} = f(\mathbf{x}_k^{(b)}, \mathbf{u}_k^{(b)})$ ;
7       Compute cost  $J^{(b)}$ ;
8       Compute gradient  $\nabla_{\mathbf{u}_{h:K-1}^{(b)}} J^{(b)}$ ;
9       Update  $\mathbf{u}_{h:K-1}^{(b)}$  with gradient descent;
10    Select top  $M$  trajectories with lowest cost and
    update  $\mu_{h:K-1}, \Sigma_{h:K-1}$ ;
11    Apply the first control  $\mathbf{u}_h$  from the best trajectory
    to the system;
12    Resample the bottom  $B - M$  trajectories from
     $\mathcal{N}(\mu_{h:K-1}, \Sigma_{h:K-1})$ ;

```

V. EXPERIMENTS AND RESULTS

A. Tactile Feedback

We evaluate our method by modeling the dynamics of the Soft Bubbles tactile sensor [30] for contact-rich robotic manipulation tasks. This sensor provides high-resolution information about deformations of a highly compliant, pressurized latex membrane. An internal camera captures changes in the membrane geometry, enabling detailed observation of contact interactions. In this work, we utilize the tactile sensor data as feedback to estimate contact forces. By combining robot proprioception with the sensor data, we obtain the pose of each hydroelastic object \mathbf{q}^{H} , partial information about the object pose \mathbf{q}^{O} , and the membrane deformations, which are used to compute the contact forces \mathbf{F}^{c} .

Specifically, to recover the contact forces \mathbf{F}^{c} from tactile observations, we use the sensed object pose and the hydroelastic object pose to determine the membrane deformation \mathbf{d}_i for each discretized point \mathbf{p}_i on the object surface. We




Fig. 4. **Simulation:** Sequential frames illustrating the four simulated manipulation tasks: planar pushing (top row), planar rotation (second row), rolling (third row), and in-hand rotation (bottom row). Note that both the bubble (pink) are modeled as hydroelastic elements and the object (solid blue) is treated as rigid. Contact forces are visualized in yellow, and goals are shown with transparency (reduced alpha).




Fig. 5. **Wrench Transmission Evaluation Setup:** The robot equipped with F/T and Soft Bubble sensors pressing a set of geometries.

then apply the compliance model described in Section IV-B to compute the forces as follows:

$$\mathbf{d}_{n,i} = \langle \mathbf{d}_i, \hat{\mathbf{n}} \rangle \hat{\mathbf{n}}, \quad (22)$$

$$\mathbf{d}_{t,i} = \mathbf{d}_i - \mathbf{d}_{n,i}, \quad (23)$$

$$\mathbf{f}_i^c = EA_i \mathbf{d}_{n,i} + K A \mathbf{d}_{t,i}. \quad (24)$$

While this approach relies on a linear approximation of the sensor compliance, it effectively corrects for dynamical mismatches and accounts for sliding drift mismatch. Also, this regularizes the smoothed dynamics and ensures a more accurate representation of the contact forces during manipulation tasks.

B. Hydroelastic Baselines

To evaluate the performance of our proposed non-holonomic (NH) model, we compare against established holonomic hydroelastic models. Specifically, we evaluate two baselines: (i) the *Pressure Field* (PF) model [8], which estimates contact interactions by integrating the full pressure distribution over the contact patch, including both normal and tangential components of force; and (ii) the *Pressure Field Frictionless* (PFF) model [8], an ablated variant of PF in which the tangential components are set to zero, thereby restricting the model to purely normal force transmission. These baselines enable us to isolate the contribution of tangential shear forces and highlight the advantages of our non-holonomic formulation in capturing realistic force transmission during contact-rich manipulation.

C. Wrench Transmission Benchmark

We evaluate the accuracy of predicted force transmission from our model against baseline methods across a diverse set of real-world contact interactions. Figure 5 visualizes the data-collection setup, where the robot is equipped with the Soft Bubble tactile sensor and makes contact with a variety of object geometries. We assess how well each model predicts the total transmitted wrench, which is measured by a Gamma ATI force-torque (F/T) sensor mounted on the robot wrist.

Table III reports the root-mean-squared error (RMSE) of the predicted wrench relative to the ground-truth F/T measurements. Our results indicate that the non-holonomic model provides the most accurate characterization of the transmitted wrench, outperforming all baselines. Interestingly, the PFF baseline — despite its simplicity compared



Fig. 6. **Real-world experiments.** Sequential frames illustrating the four manipulation tasks: planar pushing (top row), planar rotation (second row), rolling (third row), and in-hand rotation (bottom row). The orange outline indicates the object pose, and the red outline denotes the goal pose.

Object	Wrench RMSE [N] (\downarrow)					
	NH (Ours)		PF (Baseline)		PFF (Baseline)	
	Mean	Std	Mean	Std	Mean	Std
Plane	0.43	0.12	2.32	0.66	1.25	0.68
Triangle	0.44	0.22	1.31	0.40	0.97	0.38
Trapezoid	0.80	0.28	1.65	0.71	1.64	0.55
Cube	0.57	0.12	0.87	0.20	0.69	0.23
Cylinder R15	1.80	0.50	2.64	0.72	1.58	0.44
Cylinder R25	1.69	0.37	3.20	0.85	1.75	0.57

TABLE II

HYDROELASTIC MODEL BENCHMARK: WE COLLECT 40

TRAJECTORIES PER OBJECT, EACH BEING UP TO 40 STEPS. WE COMPUTE THE ROOT MEAN SQUARED ERROR (RMSE) BETWEEN THE GROUND TRUTH WRENCHES, MEASURED AT THE ROBOT WRIST, AND THE MODEL PREDICTIONS. WE REPORT THE MEAN (HIGHLIGHTED) AND STANDARD DEVIATION.

to PF — yields better wrench predictions than PF, highlighting the limitations of tangential force modeling in the PF formulation for compliant elements.

D. Contact Patch Regulation for Manipulation

In this section, we evaluate the effectiveness of our model for contact-rich manipulation tasks, focusing on its integration with gradient-based control. We consider tasks that require active regulation of the contact patch and modulation of force transmission through shear. Specifically, we study: i) planar pushing, ii) planar rotation, iii) cylindrical rolling, and iv) bi-manual in-hand rotation. These tasks are evaluated both in quasi-dynamic simulation and on real-world robotic platforms. Figures 4 and 6 illustrate the simulation and real-world environments, respectively. For each task, the robot moves the object towards 10 randomly sampled goal

configurations, repeated per method to be evaluated. For real-world execution, we use AprilTags to obtain ground-truth object poses.

Table III summarizes the quantitative evaluation results across simulation and real-world settings. The PFF baseline consistently fails to induce meaningful object motion in planar pushing, rotation, and rolling, both in simulation and in the real world, due to its inability to generate shear forces at the contact patch. In contrast, the PF baseline succeeds in simulation but exhibits unrealistic behaviors that do not transfer to real-world settings—for instance, in planar rotation and rolling tasks, it occasionally moves the object away from the goal. Overall, our proposed method outperforms both baselines, reducing errors across all tasks. While the real-world execution error remains higher than in simulation, our approach narrows the sim-to-real gap more effectively than competing methods. As expected, closing the control loop improves performance compared to open-loop execution.

VI. DISCUSSION, LIMITATIONS, AND FUTURE WORK

The proposed non-holonomic hydroelastic model, combined with tactile feedback, shows strong potential for accurately estimating contact forces and compensating for dynamical mismatches in contact-rich manipulation. Nevertheless, several limitations remain, pointing to avenues for future work. A key limitation lies in the state representation: the explicit inclusion of contact forces causes the state dimension to scale poorly with object discretization and the number of objects. This results in substantial computational overhead, particularly for methods such as iterative Model Predictive Control (iMPC), where the computation of large Jacobian

Task	Non-holonomic (Ours)				Pressure Field (Baseline)				PF Frictionless (Baseline)			
	OL Error [mm]		CL Error [mm]		OL Error [mm]		CL Error [mm]		OL Error [mm]		CL Error [mm]	
	Mean ↓	Std ↓	Mean ↓	Std ↓	Mean ↓	Std ↓	Mean ↓	Std ↓	Mean ↓	Std ↓	Mean ↓	Std ↓
Simulation												
Planar Pushing	2.21	0.68	0.02	0.01	2.79	0.74	2.54	1.59	101	2.07	100	2.87
Planar Rotation	3.20	0.88	0.04	0.05	32.6	17.9	20.1	22.3	40.7	10.1	40.6	12.9
Rolling	2.24	1.83	6.14	5.56	41.5	24.0	42.8	26.8	47.3	18.4	49.1	18.8
In-hand Rotation	3.98	1.39	0.66	1.03	8.94	2.91	7.37	3.34	9.54	3.77	9.58	3.79
Real World												
Planar Pushing	6.59	3.62	2.09	1.02	58.5	9.63	45.3	16.7	73.3	28.4	90.6	20.2
Planar Rotation	14.9	9.34	3.87	1.23	57.5	17.3	57.8	16.5	37.9	7.25	37.9	7.28
Rolling	21.2	15.0	9.02	2.86	61.1	18.0	53.1	26.9	52.8	19.6	62.5	23.1
In-hand Rotation	5.19	0.84	4.07	0.89	12.7	4.75	9.14	4.78	14.3	8.34	10.8	4.67

TABLE III

OPTIMIZED TRAJECTORY CONTROL EVALUATION: BENCHMARK OF DIFFERENT MODELS (NON-HOLONOMIC, PF, PF FRICTIONLESS) IN BOTH SIMULATION AND REAL-WORLD EXPERIMENTS. EACH CELL REPORTS MEAN AND STD FOR OPEN-LOOP (OL) AND CLOSED-LOOP (CL) POSE TRACKING ERRORS [MM]. MEAN COLUMNS ARE HIGHLIGHTED.

matrices for locally linearized dynamics becomes prohibitive. Moreover, our method currently applies only in closed-loop scenarios where contact forces can be directly observed or reliably estimated, for example via tactile sensing. Integrating recent approaches for contact force estimation [31] [24] could help alleviate this restriction and broaden applicability. Finally, while in this work we primarily evaluate the model within a trajectory optimization setting, future research could explore its use in reinforcement learning. Given its ability to capture the dynamics of highly compliant contacts with relatively low computational overhead, the model may serve as a useful tool for training policies in simulation that transfer effectively to real-world systems.

REFERENCES

- [1] M. T. Mason, "Toward robotic manipulation," *Annual Review of Control, Robotics, and Autonomous Systems*, 2018.
- [2] N. Hogan, "Impedance control: An approach to manipulation," in *1984 American control conference*, 1984.
- [3] G. A. Pratt and M. M. Williamson, "Series elastic actuators," *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, 1995.
- [4] M. Oller, M. P. i Lisboa, D. Berenson, and N. Fazeli, "Manipulation via membranes: High-resolution and highly deformable tactile sensing and control," in *Conference on Robot Learning*, PMLR, 2022.
- [5] H. T. Suh, N. Kuppusswamy, T. Pang, P. Mitiguy, A. Alspach, and R. Tedrake, "Seed: Series elastic end effectors in 6d for visuotactile tool use," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [6] M. Oller, D. Berenson, and N. Fazeli, "Tactile-driven non-prehensile object manipulation via extrinsic contact mode control," *arXiv preprint arXiv:2405.18214*, 2024.
- [7] J.-C. Peng, S. Yao, and K. Hauser, "3d force and contact estimation for a soft-bubble visuotactile sensor using fem," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [8] 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 International Conference on Intelligent Robots and Systems (IROS)*.
- [9] J. Masterjohn, D. Guoy, J. Shepherd, and A. Castro, "Velocity level approximation of pressure field contact patches," *IEEE Robotics and Automation Letters*, 2022.
- [10] M. Anitescu and F. A. Potra, "Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems," *Nonlinear Dynamics*, 1997.
- [11] D. Stewart and J. C. Trinkle, "An implicit time-stepping scheme for rigid body dynamics with coulomb friction," in *IEEE International Conference on Robotics and Automation*, IEEE, 2000.
- [12] M. Anitescu, "Optimization-based simulation of nonsmooth rigid multibody dynamics," *Mathematical Programming*, 2006.
- [13] T. Pang and R. Tedrake, "A convex quasistatic time-stepping scheme for rigid multibody systems with contact and friction," in *International Conference on Robotics and Automation (ICRA)*, IEEE, 2021.
- [14] E. Todorov, "Implicit nonlinear complementarity: A new approach to contact dynamics," in *2010 IEEE ICRA*.
- [15] W. Jin, "Complementarity-free multi-contact modeling and optimization for dexterous manipulation," *arXiv:2408.07855*, 2024.
- [16] H. J. T. Suh, T. Pang, and R. Tedrake, "Bundled gradients through contact via randomized smoothing," *IEEE Robotics and Automation Letters*, 2022.
- [17] T. Pang, H. T. Suh, L. Yang, and R. Tedrake, "Global planning for contact-rich manipulation via local smoothing of quasi-dynamic contact models," *IEEE Transactions on robotics*, 2023.
- [18] R. D. Howe and M. R. Cutkosky, "Practical force-motion models for sliding manipulation," *IJRR*, 1996.
- [19] S. Goyal, A. Ruina, and J. Papadopoulos, "Limit surface and moment function descriptions of planar sliding," in *1989 IEEE International Conference on Robotics and Automation*.
- [20] X. Yi, A. Dang, and N. Fazeli, "Dual asymmetric limit surfaces and their applications to planar manipulation," *Autonomous Robots*, 2024.
- [21] M. M. Ghazaei Ardakani, J. Bimbo, and D. Prattichizzo, "Quasi-static analysis of planar sliding using friction patches," *IJRR*, 2020.
- [22] Y. Hou, Z. Jia, and M. T. Mason, "Manipulation with shared grasping," *arXiv:2006.02996*, 2020.
- [23] N. Kuppusswamy, A. Castro, C. Phillips-Grafflin, A. Alspach, and R. Tedrake, "Fast model-based contact patch and pose estimation for highly deformable dense-geometry tactile sensors," *IEEE Robotics and Automation Letters*, 2019.
- [24] Y. Wi, J. Lee, M. Oller, and N. Fazeli, "Neural inverse source problems," *arXiv preprint arXiv:2411.01665*, 2024.
- [25] V. Kurtz and H. Lin, "Contact-implicit trajectory optimization with hydroelastic contact and ilqr," in *International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2022.
- [26] J. Xu, S. Kim, T. Chen, A. R. Garcia, P. Agrawal, W. Matusik, and S. Sueda, "Efficient tactile simulation with differentiability for robotic manipulation," in *Conference on Robot Learning*, 2023.
- [27] I. Akinola, J. Xu, J. Carius, D. Fox, and Y. Narang, "TacsI: A library for visuotactile sensor simulation and learning," *IEEE TRO*, 2025.
- [28] B. Aceituno-Cabezas and A. Rodriguez, "A global quasi-dynamic model for contact-trajectory optimization in manipulation," 2020.
- [29] H. Bharadhwaj, K. Xie, and F. Shkurti, "Model-predictive control via cross-entropy and gradient-based optimization," in *Learning for Dynamics and Control*, PMLR, 2020.
- [30] A. Alspach, K. Hashimoto, N. Kuppusswamy, and R. Tedrake, "Soft-bubble: A highly compliant dense geometry tactile sensor for robot manipulation," in *RoboSoft*, IEEE, 2019.
- [31] M. Van der Merwe, M. Oller, D. Berenson, and N. Fazeli, "Estimating deformable-rigid contact interactions for a deformable tool via learning and model-based optimization," *IEEE Robotics and Automation Letters*, 2025.