

2024

LeTac-MPC: Learning Model Predictive Control for Tactile-Reactive Grasping

Zhengtong Xu

Yu She

Follow this and additional works at: <https://docs.lib.purdue.edu/iepubs>



Part of the [Industrial Engineering Commons](#)

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries.
Please contact epubs@purdue.edu for additional information.

LeTac-MPC: Learning Model Predictive Control for Tactile-reactive Grasping

Zhengtong Xu, Yu She*




Fig. 1: LeTac-MPC is capable of reactive grasping for objects with varying physical properties, shapes, sizes, and surface textures. It enables robust grasping in dynamic manipulation scenarios, adjusting the gripper width to provide the appropriate grasping force based on tactile feedback. This ensures stability without damaging the grasped object. Our demonstrations show the robot performing various tasks using LeTac-MPC. Left: The robot grasps a deformable object. Middle and right: LeTac-MPC allows the gripper to maintain a stable grasping under dynamic shaking and unpredictable external collisions, respectively.

Abstract—Grasping is a crucial task in robotics, necessitating tactile feedback and reactive grasping adjustments for robust grasping of objects under various conditions and with differing physical properties. In this paper, we introduce LeTac-MPC, a learning-based model predictive control (MPC) for tactile-reactive grasping. Our approach enables the gripper to grasp objects with different physical properties on dynamic and force-interactive tasks. We utilize a vision-based tactile sensor, GelSight [1], which is capable of perceiving high-resolution tactile feedback that contains information on the physical properties and states of the grasped object. LeTac-MPC incorporates a differentiable MPC layer designed to model the embeddings extracted by a neural network (NN) from tactile feedback. This design facilitates convergent and robust grasping control at a frequency of 25 Hz. We propose a fully automated data collection pipeline and collect a dataset only using standardized blocks with different physical properties. However, our trained controller can generalize to daily objects with different sizes, shapes, materials, and textures. The experimental results demonstrate the effectiveness and robustness of the proposed approach. We compare LeTac-MPC with two purely model-based tactile-reactive controllers (MPC and PD) and open-loop grasping. Our results show that LeTac-MPC has optimal performance in dynamic and force-interactive tasks and optimal generalizability. We release our code and dataset at <https://github.com/ZhengtongXu/LeTac-MPC>.

Index Terms—Tactile control, deep learning in robotics and automation, perception for grasping and manipulation.

*Address all correspondence to this author.

Zhengtong Xu and Yu She are with the School of Industrial Engineering, Purdue University, West Lafayette, USA (E-mail: {xu1703, shey}@purdue.edu).

I. INTRODUCTION

Grasping is a reactive task in which humans estimate the states and physical properties of the grasped object by tactile feedback of the fingers and dynamically adjust finger behavior. This allows for robust grasping regardless of the object’s stiffness, whether it needs to be stationary or dynamically moved, and whether it is subjected to external forces. An ideal grasping algorithm for robots should be able to do the same. In this paper, we study the task of robotic tactile-reactive grasping. We specifically focus on a vision-based tactile sensor, GelSight [1], due to its advantages of perceiving high-resolution tactile feedback that contains rich information on the physical properties and states of the grasped object.

In recent years, many works demonstrate that tactile sensors can perceive various information of objects, such as texture [2]–[4], material [5], [6], shape [7], [8], and hardness [9]. However, utilizing tactile feedback to design robotic tactile-reactive grasping controllers faces several challenges.

1) **Complex object physical properties:** Objects may have different physical properties, which poses great challenges to tactile-based perception. Existing tactile perception methods often assume that grasped objects are rigid [10], [11], and sensors, such as vision-based tactile sensors [12]–[15], may not be sensitive if the grasped object is softer than the sensor elastomer [1]. As a result, it can be difficult for controllers to get high-quality and stable feedback signals using model-based extraction methods when handling objects

with different physical properties, and it can be challenging to generalize the same set of controller parameters to different objects, as their physical properties and dynamics may differ.

2) **Dynamic and force-interactive tasks:** Statically grasping an object is often insufficient for dynamic and force-interactive tasks. For example, when robots use the gripper to pick fruit, there is a force interaction between the gripper, the fruit, and the branches. Similarly, in dynamic pick-and-place tasks, the grasped object may experience inertia force and collide with unexpected obstacles during transportation. In both of these scenarios, if a gripper is not responsive enough to state changes, external disturbances, and force interactions, it can result in the grasped object falling from the gripper. In addition, if a larger grasping force is used for a stronger grasp, delicate objects such as fruits, eggs, bread, and crackers can be destroyed. Therefore, it is challenging to design a reactive controller that can grasp with appropriate force according to the physical properties and states of the object.

3) **High-resolution tactile feedback integration:** Another challenge is the integration of high-resolution tactile feedback, such as the vision-based tactile sensor GelSight [1], which contains rich information, into a real-time control loop for tactile-reactive grasping. Traditional control methods are based on low-dimensional feedback signals, which require the extraction of required information from tactile images for reduction of dimensionality [15]–[17]. However, many feature extraction methods make assumptions about the physical properties and shape of the object [10], [11], [16], making it difficult to generalize to objects with different physical properties and shapes. On the other hand, robot learning methods are based on high-dimensional observations such as visual and tactile images. However, these methods often overlook crucial aspects of the control process, such as convergence, response speed, control frequency, and constraints. As a result, they are more suitable for generating open-loop or low-frequency actions, but not well-suited for tasks that demand rapid and reactive responses.

To address these challenges, we propose LeTac-MPC, a learning-based model predictive control (MPC) for tactile-reactive grasping. The novelties of the proposed approach are as follows.

1) We design a neural network (NN) architecture with a differentiable MPC layer as the output layer. The MPC layer is modeled according to the control objectives to rationalize the embeddings extracted by the NN from the tactile feedback. The combination of NN and MPC layer provides the advantages of strong generalizability and ease of implementation as a real-time controller. We implement the trained NN and MPC layer in a model predictive controller that is convergent and can run at a frequency of 25 Hz. This controller enables robust reactive grasping of various daily objects in dynamic and force-interactive tasks.

2) To train the model, we use four standardized blocks of different materials with different physical properties to collect a dataset and propose a fully automated data collection pipeline. Finally, we show that despite our data collection

with standardized blocks in terms of sizes, shapes, and materials, and with no textures on the block surface, our trained controller can generalize to various daily objects with different sizes, shapes, materials, and textures. This shows the strong generalizability of LeTac-MPC.

3) We also compare two purely model-based tactile-reactive controllers (MPC and PD) as baseline methods. Through experimental comparisons between these two model-based controllers, LeTac-MPC, and open-loop grasping, we show that LeTac-MPC has optimal performance in dynamic and force-interactive tasks and optimal generalizability.

II. RELATED WORK

This section provides a summary of previous research related to exploiting tactile sensing for control and learning.

A. Tactile Control

Previous research has shown that tactile feedback provided by vision-based tactile sensors can be used to design manipulation controllers for rigid objects [11], [17], [18], cables [16], and clothes [19]. These methods extract state estimations of the manipulated object, such as pose and contact line with the environment, from tactile feedback, and then design controllers. However, these approaches are not generalizable to objects with different physical properties due to limitations in their model-based state extraction methods. The work presented in [20] proposes a tactile control framework for robotic pushing that uses tactile feedback to estimate and predict state information. In [21], a learning-based method is proposed for flipping pages via tactile sensing. However, these two works do not fully use the rich information about the object's physical properties that can be obtained from tactile feedback, which is critical for enhancing generalization to different objects.

Recently, the work in [22] proposes a method to learn soft tactile sensor membrane dynamics. However, this method does not consider the dynamics of the manipulated object, which is important for manipulating deformable objects. The work in [23] proposes a learning-based tactile MPC framework. However, this framework is not suitable for real-time control tasks because it uses a video prediction model, resulting in high computational cost. As mentioned in the paper, during the experimental testing, this MPC only runs at around 1 Hz.

B. Tactile Grasping

Multiple works exploit tactile feedback and NN to predict grasp success and stability [24]–[27]. Based on the predictive ability of NN, the works in [28]–[31] sample potential grasping actions and select the action with the lowest cost as the optimal grasping. However, these methods do not focus primarily on real-time control. They do not consider important aspects such as convergence, response speed, control frequency, and constraints in the control process. Consequently, they are better suited for low-frequency action

generation rather than real-time control. For example, in [30], the computation times of different NN models for one trial are both more than 0.28 s. As a result, these methods cannot respond fast and smooth enough when the robot performs dynamic and force-interactive tasks. The works in [10], [32] maintain stable grasping by slip detection. Therefore, the accuracy of slip detection determines the performance of the method. However, in practical situations, slip detection may be insufficiently sensitive due to the different shapes and physical properties of grasped objects.

In [16], [33], grasping controllers with model-based tactile feature extraction are proposed. However, these methods cannot be widely applicable to objects with different shapes and physical properties since some feature extraction methods only work on specific types of object. For example, GelSight, as used in [16], [33], cannot obtain obvious features for objects softer than the sensor gel [1]. These features are extremely difficult to process and extract using model-based methods. Instead, existing work on tactile perception shows that NN can effectively extract these features [9]. Therefore, it is crucial to expand learning-based tactile perception to develop a learning-based tactile control method that can generalize to objects with different shapes, textures, and physical properties.

C. Learning Physical Properties for Manipulation

Learning-based methods that incorporate tactile and visual feedback can achieve dynamic manipulation without prior knowledge of the physical properties of the manipulated object [34]–[36]. However, these methods are not suitable for tasks that require reactive behavior. In contrast, many manipulation tasks require real-time sensing and real-time control. In [37], [38], the authors propose methods that utilize neural radiance fields (NeRF) for manipulating objects with complex properties. However, these methods are not suitable for tasks demanding real-time control and fast response due to the relatively high computational costs. In [39], a model learning method for elastic deformable linear objects is proposed that enables real-time cable manipulation. The input of its model is the cable state extracted from the raw image, which requires markers on the cable for tracking. However, in practice, extracting cable states from images without markers is challenging. In summary, achieving real-time, learning-based robot control directly using high-dimensional images as input remains a significant challenge.

The work in [40] proposes a learning-based real-time tactile control method to swing up rigid poles with different physical properties. However, there is no research on real-time tactile controllers for soft and deformable objects.

III. APPROACH

This section outlines our proposed learning-based MPC for tactile-reactive grasping. Initially, we design an NN architecture that includes a differentiable MPC layer (Sections III-A and III-B) to extract and represent the diverse and complex physical properties of grasped objects. Then,

we introduce the details of training our proposed model (Section III-C). Furthermore, we describe our automated data collection pipeline in Section III-D. Finally, we present the implementation of the NN as a real-time controller in Section III-F.

By our method, the gripper can successfully grasp objects with different physical properties and shapes while applying the appropriate force to avoid damage to fragile and delicate items. Our method can also adjust the grasping in real-time based on tactile feedback to improve the grasping robustness when performing dynamic and force-interactive tasks.

A. Tactile Information Encoding

We use GelSight [1], a vision-based tactile sensor that provides high-resolution images of the contact surface geometry and strain field, as shown in Fig. 1. To extract states and physical properties of grasped objects, we use a convolutional neural network (CNN) with a multi-layer perceptron (MLP) to encode a tactile image into a low-dimensional embedding $\mathbf{f} \in \mathbb{R}^M$, where M is the dimension of the embedding. The works on tactile perception in [9], [41] demonstrate effective feature extraction from tactile images using pre-trained visual models. In our implementation, we use a pre-trained ResNet-152 [42] as the CNN architecture and replace the last layer with a two-layer MLP with ReLU activation.

B. Differentiable MPC Layer

In this section, we introduce a differentiable optimization layer that utilizes the extracted embeddings for controller design, as depicted in Fig. 2. This layer formulates an MPC problem, which we henceforth refer to as a differentiable MPC layer.

Differentiable optimization layers [43] function similarly to standard layers in neural networks, capable of performing both forward passes and backpropagation in batch form. During the forward pass, the layer takes in variables that define an optimization problem and outputs results related to the solution of that problem. Consequently, the forward pass of this layer involves solving optimization problems in batch form. Backpropagation in the differentiable optimization layer involves updating the parameters of the optimization problem based on the calculated gradients.

In this paper, we consider the two-finger gripper with one degree of freedom linear motion. This kind of gripper is widely used because of its practicality. Denoting the sampling time interval as Δt , we can first formulate the gripper motion model as:

$$\begin{bmatrix} p_{n+1} \\ v_{n+1} \end{bmatrix} = \mathbf{A}_g \begin{bmatrix} p_n \\ v_n \end{bmatrix} + \mathbf{B}_g a_n, \quad (1)$$

$$\text{where } \mathbf{A}_g = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \in \mathbb{R}^{2 \times 2}, \mathbf{B}_g = \begin{bmatrix} \frac{1}{2} \Delta t^2 \\ \Delta t \end{bmatrix} \in \mathbb{R}^2.$$

Scalars p , v , and a are the position, velocity, and acceleration of the motion of gripper fingers. Note that p , v , and a do not refer to the motion of the gripper itself in Cartesian space




Fig. 2: LeTac-MPC network model. We use the raw image of tactile feedback as the model input. ℓ is our proposed loss function.

but rather denote the one-degree-of-freedom motion of the gripper fingers. The right subscript n represents the n -th time step.

In order to utilize the embeddings $\mathbf{f} \in \mathbb{R}^M$ from a NN to design a controller, we propose a differentiable MPC layer that takes into account the fact that the features of tactile images change as the gripper width changes when grasping an object. We formulate the relationship between the tactile embeddings and gripper states as a linear model

$$\mathbf{f}_{n+1} = \mathbf{f}_n + \mathbf{A}_f v_n,$$

where $\mathbf{A}_f \in \mathbb{R}^M$ would be learned. This modeling approach does not lose its applicability to complex dynamics because the tactile embeddings are actually generated by an NN, which has strong nonlinear fitting capabilities. Moreover, for MPC, we do not expect the model to maintain high accuracy over long timescales because MPC operates at a high frequency with receding horizon control. Therefore, as long as the model demonstrates adequate representational capability in a short-term local area, it is sufficient for MPC performance. The combination of the NN and the MPC layer provides the advantages of strong generalizability and ease of implementation as a real-time controller. The forward pass of the MPC layer involves solving the following optimization problem:

$$\begin{aligned} \mathbf{a}_n^* = \arg \min_{\mathbf{a}_n} & P(\mathbf{f}_{n+N}^T \mathbf{Q}_f \mathbf{f}_{n+N} + Q_v v_{n+N}^2) \\ & + \sum_{k=n}^{n+N-1} \mathbf{f}_k^T \mathbf{Q}_f \mathbf{f}_k + Q_v v_k^2 + Q_a a_k^2, \end{aligned} \quad (2)$$

subject to

$$\begin{bmatrix} \mathbf{f}_{n+1} \\ p_{n+1} \\ v_{n+1} \end{bmatrix} = \begin{bmatrix} \mathbf{I}^{M \times M} & \bar{\mathbf{A}}_f \\ \mathbf{0}^{2 \times M} & \mathbf{A}_g \end{bmatrix} \begin{bmatrix} \mathbf{f}_n \\ p_n \\ v_n \end{bmatrix} + \begin{bmatrix} \mathbf{0}^{M \times 1} \\ \mathbf{B}_g \end{bmatrix} a_n, \quad (3)$$

$$\text{where } \bar{\mathbf{A}}_f = \begin{bmatrix} \mathbf{0}^{M \times 1} & \mathbf{A}_f \end{bmatrix} \in \mathbb{R}^{M \times 2}.$$

n denotes the current time step, and we can define the initial time step as $n = 0$. As time progresses, there is no upper limit to n . Therefore, $n = 0, 1, 2, 3, \dots$. $\mathbf{a}_n = [a_n, a_{n+1}, \dots, a_{n+N-1}]^T \in \mathbb{R}^N$ is the acceleration sequence, N is the prediction horizon, the vector $\mathbf{A}_f \in \mathbb{R}^M$ is part of the state transition matrix, $\mathbf{Q}_f \in \mathbb{R}^{M \times M}$ is the weight matrix for $\mathbf{f}_n$, and Q_v and Q_a are the weight coefficients for

v_n and a_n , respectively. The scalar P is used to amplify the terminal cost to speed up convergence. Note that the “ n ” in v_n, v_{n+1} represents the current time step n , while the “ N ” in v_{n+N} represents the MPC horizon length N . Importantly, we design $\mathbf{f}_{n+1}$ is only directly related to $\mathbf{f}_n$ and v_n . This design helps the model avoid the influence of the different sizes of the grasped objects. Therefore, the trained model can be generalized to objects with different sizes naturally.

Upon solving the optimization problem (2) to obtain $\mathbf{a}_n^* = [a_n^*, a_{n+1}^*, \dots, a_{n+N-1}^*]^T \in \mathbb{R}^N$, we can generate the entire trajectory of the gripper motion output by the MPC controller, including p^*, v^* and a^* , using equation (1). Ultimately, the final output of the MPC layer consists of the position sequence $[p_{n+1}^*, \dots, p_{n+N}^*]^T \in \mathbb{R}^N$, as shown in Fig. 2. The rationale behind this design is that most grippers, such as the WSG gripper used in this paper, have a low-level tracking controller capable of effectively tracking high-level position commands. By outputting the position sequence of the trajectory generated by the MPC layer and sending it to the low-level controller, we can approximate the tracking of the entire trajectory, including p^*, v^* and a^* , in a simple yet effective manner. For more details, see the experiments described in Section V.

C. Training

For this MPC layer, we choose $\mathbf{A}_f \in \mathbb{R}^M$ and $\mathbf{Q}_f \in \mathbb{R}^{M \times M}$ as the parameters we want to learn, since they cannot be derived from the model-based formulation. On the other hand, the scalars Q_v , Q_a , and P are model-based control parameters that are directly related to the response and convergence speed of the MPC. Therefore, we do not consider them as parameters to learn and can assign reasonable values directly. Additionally, it is important to note that Q_v , Q_a , and P must be positive. The positivity of these scalars is crucial for the feasibility of optimization problem (2). We will explain why these scalars must be positive in Theorem 1.

Optimization problem (2) is a quadratic program (QP). The work in [43] demonstrates that a QP can be represented as a differentiable layer. In addition, a method is proposed to solve multiple QPs in batch form in [43]. Therefore, the NN shown in Fig. 2 can perform forward pass and backpropagation in batch form if the optimization problem (2) is feasible throughout the training process. To ensure feasibility, the resulting QP’s Hessian matrix must remain

symmetric positive definite when $\mathbf{A}_f$ and $\mathbf{Q}_f$ change during training [44].

Regarding the feasibility of our proposed differentiable MPC layer, we state the following theorem.

Theorem 1. *If $Q_v, Q_a, P > 0$ and $\mathbf{Q}_f$ is symmetric positive definite, then the resulting QP from equations (2) and (3) is feasible regardless of any changes to the dimension of the embedding M , the prediction horizon N , and $\mathbf{A}_f$.*

Proof. Since the optimization problem (2) does not have constraints besides (3), the resulting QP is unconstrained. Therefore, to prove the theorem, it is sufficient to demonstrate that the Hessian matrix of the resulting QP is symmetric positive definite [44]. The resulting Hessian matrix from (2) is

$$\mathbf{H} = 2(\bar{\mathbf{Q}}_a + \bar{\mathbf{S}}^T \bar{\mathbf{Q}} \bar{\mathbf{S}}), \quad (4)$$

where $\bar{\mathbf{Q}}_a = \text{blkdiag}(Q_a, \dots, Q_a) \in \mathbb{R}^{N \times N}$,

$$\bar{\mathbf{Q}} = \text{blkdiag}(\mathbf{Q}, \dots, \mathbf{Q}, P\mathbf{Q}) \in \mathbb{R}^{N(M+1) \times N(M+1)},$$

$$\mathbf{Q} = \text{blkdiag}(\mathbf{Q}_f, Q_v) \in \mathbb{R}^{(M+1) \times (M+1)},$$

$$\bar{\mathbf{S}} = \begin{bmatrix} \mathbf{B} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{A}\mathbf{B} & \mathbf{B} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}^{N-1}\mathbf{B} & \mathbf{A}^{N-2}\mathbf{B} & \dots & \mathbf{B} \end{bmatrix}, \quad (5)$$

$$\bar{\mathbf{S}} \in \mathbb{R}^{N(M+1) \times N},$$

$$\mathbf{A} = \begin{bmatrix} \mathbf{I}^{M \times M} & \mathbf{A}_f \\ \mathbf{0}^{1 \times M} & 1 \end{bmatrix} \in \mathbb{R}^{(M+1) \times (M+1)},$$

$$\mathbf{B} = \begin{bmatrix} \mathbf{0}^{M \times 1} \\ \Delta t \end{bmatrix} \in \mathbb{R}^{M+1}.$$

Because $Q_v, Q_a, P > 0$ and $\mathbf{Q}_f$ are symmetric positive definite, $\bar{\mathbf{Q}}$ and $\bar{\mathbf{Q}}_a$ are symmetric positive definite. Therefore, by equation (4), we can see that the Hessian matrix is symmetric positive semi-definite.

To prove that the Hessian matrix is symmetric positive definite, we can prove $\text{rank}(\bar{\mathbf{S}}) = N$. Because $\mathbf{B} \neq \mathbf{0}$, we can clearly see from equation (5) that $\text{rank}(\bar{\mathbf{S}}) = N$. Therefore, we have shown that the Hessian matrix of the resulting QP is symmetric positive definite, which completes the proof of the theorem. $\square$

Theorem 1 demonstrates that if we ensure that $\mathbf{Q}_f$ is symmetric positive definite, the selection of N and M , and uncertainty in $\mathbf{f}_n$ and $\mathbf{A}_f$ will not affect the QP's feasibility. As a result, we can freely tune the parameters and train the model stably. To ensure that $\mathbf{Q}_f$ remains symmetric positive definite during training, we utilize a Cholesky factorization

$$\mathbf{Q}_f = \mathbf{L}_f \mathbf{L}_f^T + \epsilon \mathbf{I}^{M \times M},$$

and directly learn $\mathbf{L}_f$, where $\mathbf{L}_f$ is a lower triangular matrix and ϵ is a very small scalar. We choose $\epsilon = 1 \times 10^{-4}$.

Although Theorem 1 shows that the selection of the prediction horizon N and the embedding dimension M does not affect the feasibility of the QP, it is still important to choose the appropriate values for these parameters. Very small N and M can negatively impact the network's ability to extract

useful information and generalize well. Conversely, setting N and M too large can slow down the optimization problem's solution due to the increase in dimensionality when implementing the trained model as a controller, which can result in a low control frequency. Moreover, when N is excessively large, the linear MPC layer may result in a large prediction error. Therefore, the choice of N and M should balance the expressiveness of the model and the computational efficiency of the controller. In our implementation, we choose $N = 15$ and $M = 20$ to strike this balance.

Denote the loss function as ℓ (see Section III-E), "differentiable" in our formulation means, gradients of ℓ with respect to $\mathbf{A}_f$ and $\mathbf{L}_f$, denoted as $\nabla_{\mathbf{A}_f}^\ell$ and $\nabla_{\mathbf{L}_f}^\ell$, can be computed with standard backpropagation. Additionally, by employing the batched QP solver introduced in [43], our proposed differentiable MPC layer can be integrated into a NN, enabling the training of the entire model in a standardized manner.

D. Data Collection

We selected four different types of materials for data collection, as illustrated in Fig. 3. It is important to mention that we chose only standardized blocks without surface texture and do not require specific sizes, as long as they fit the gripper. By doing so, the data collection process is simplified. we show in Section V that our proposed network model can learn from these standardized blocks and generalize to other daily objects with different textures, sizes, shapes, and physical properties.

These blocks are used as grasped objects to capture raw tactile images. The tactile feedback of these blocks varies depending on their unique physical properties. As shown in Fig. 3, we can see the differences in tactile feedback between these 4 blocks.

The contact area of hard materials is smaller but more distinct, while the contact area of soft materials is larger but subtler. This is because of the difference of the objects' stiffness. When objects with different stiffness are grasped under the same applied force, the deformation of the grasped objects and the resulting deformation on the GelSight surface will be different, resulting in different raw tactile images.

The 3D printing block serves the role of a rigid object. We utilized a Markforged Onyx printer, which prints using micro carbon fiber filled nylon. This material is Onyx composite material. Its tensile modulus is 2.4 GPa, and its flexural modulus is 3 GPa [45]. When printed with a block form factor, it behaves rigid when grasped by our grippers. Compared with the hard block made by 3D printing, hard rubber is softer, so the contact area is larger. Additionally, the contact area of gel is the most subtle. The gel is made of XP-565 with shore A hardness of 27 [46], which is approximate 0.0009942 GPa of the Young's modulus [47]. The stiffness of the softer rubber is between the gel and hard rubber. Finally, the difference in marker displacements can be attributed to the fact that materials with higher surface friction coefficients tend to induce a more pronounced displacement on the GelSight's surface. It is worth noting that XP-565 is a commonly used




Fig. 3: For each type of material, bottom left is the raw tactile image overlapping with the marker tracking when grasping the corresponding block and bottom right is the raw tactile image without marker tracking. These tactile images are collected under the same grasping force. Our goal here is to show that different materials have different physical properties, which leads to different features in the tactile image under the same grasping conditions. Therefore, we do not estimate the specific values of tangential force here.




Fig. 4: Left: data collection setup. Right: automated data collection pipeline. The xyz -coordinate of the end-effector's position increments is shown in the left figure.




Fig. 5: Collected data visualization.

material to make silicone gel for GelSight sensors. The gel made of XP-565 has a nominal shore A hardness of 27 [46]. Therefore, the stiffness of the gel in Fig. 3 is very close to the gel of the GelSight sensors made by XP-565. We believe that this point can serve as a benchmark for selecting data collection blocks: the chosen blocks usually possess a stiffness greater than that of the sensor gel. Otherwise, the

deformation caused by the block on the sensor gel would be very subtle or even be invisible.

We believe that when other people select materials to construct a dataset, following our approach of choosing one hard material (such as rigid 3D printed objects or wood) and then selecting materials with progressively lower stiffness, up to a total of four materials, would be an effective process for dataset construction.

Our network input consists only of raw images. Since marker displacements are actually obtained by optical flow from raw images, the raw images themselves already contain this information. To collect data for training the NN shown in Fig. 2, we propose an automatic data collection pipeline utilizing a dual-arm setup. We use Franka Panda robotic arm. One of the arms serves as the leader arm, and we control it by making it follow specific trajectories. The other arm serves as the follower arm, and we implement an impedance controller for it. The WSG 50-110 grippers are attached to the end-effectors of both arms, with one finger on the follower arm mounting the GelSight tactile sensor.

A complete data collection trial is illustrated in Fig. 4. We perform multiple iterations of this trial to collect the necessary data. At the start of each trial, both grippers securely grasp the object. We randomly select position increments Δx and Δy within the range of $|\Delta x| < 35$ mm and $|\Delta y| < 21$ mm, and apply these increments to the end-

effector of the leader arm. The leader arm then moves, and the follower arm follows its motion due to the impedance controller. As a result, an external force is horizontally applied to the grasped object, with the magnitude and direction determined by the deviation Δx and Δy of the follower manipulator's end-effector from the equilibrium point of its impedance controller.

Next, the follower gripper gradually relaxes with a velocity of 4.5 mm/s, and we record the width of the follower gripper p_n and the raw image of the tactile feedback at 60 Hz. When the follower gripper's grasping force applied to the object cannot provide a maximum static friction force greater than the external force from the impedance controller, slipping occurs between the follower gripper and the object. This results in the follower arm moving back to the equilibrium point of the impedance controller. We record the width of the follower gripper when slipping occurs as the label for all raw images recorded during the entire trial, denoted as p_{slip} .

Ideally, we hope that the leader gripper maintains an absolutely tight grasping throughout the process and there is no slipping between the leader gripper and the grasped object. To achieve that, we maximize the grasping force of the leader gripper on the data collection block and enhance the friction between the leader gripper's fingers and the block by adding a gel layer on the finger. Naturally, we could not guarantee that there would be no movement between the object and the leader gripper during extended periods of data collection. However, our setup ensures that any movement, if it occurred, is minimal and does not affect data collection. Therefore, when the leader gripper stays static, the grasped object stays static and we set the velocity of the end-effector of the follower arm to be greater than 5 mm/s as the condition for indicating slipping. Once slipping occurs for 0.2 s, we stop recording the raw images. The reasons of selecting 0.2 s are as follows. For each material, as shown in Fig. 5, when the external force varies from 0 to 12N, the fluctuation range of the p_{slip} collected for each material is approximately within 2 mm. Using 2 mm as a reference, to ensure the data collected in a reasonable range (for example, to avoid collecting too many images without contact), we chose a distance of 0.9 mm ($0.2 \text{ s} \times 4.5 \text{ mm/s}$) as the gripper width range of one data collection trial. It is worth noting that the choice of 0.2 s here is approximate. The goal is to prevent the collected data from becoming overly biased.

For the gripper velocity v_n in the dataset, we randomly select a value within the range of $(-1 \text{ mm/s}, 1 \text{ mm/s})$ for each data point. As shown in Fig. 2, p_n, v_n , and the raw tactile image are the inputs of our model.

Finally, the follower gripper continues to relax until the object is completely released. After that, the leader arm returns to its initial position, and the follower gripper tightly re-grasps the object. The next trial then begins.

The Franka Panda robot features force-torque feedback of the end-effector, eliminating the need for additional sensors. We simultaneously record force feedback during the data collection process. Although force feedback is not useful for model training, we can visualize its relationship with

p_{slip} to verify the reasonableness of the collected data, as shown in Fig. 5. Fig. 5 shows that the gripper width and the external force applied to the block are approximately negatively correlated. This indicates an approximate linear positive correlation between the pressure and the total external force applied to each block, aligned with the Coulomb friction model. According to this model, the maximum static friction force $F_{\text{max}} = \mu F_{\text{normal}}$, where F_{normal} is the normal force exerted on the contact surface and μ is the coefficient of static friction. Additionally, we know that the gripper width and F_{normal} are negatively correlated (a smaller gripper width means tighter grasping and a larger F_{normal}). Therefore, the depiction in Fig. 5 corroborates this equation. This demonstrates the effectiveness of our automated data collection method.

We also record some raw tactile images when not applying external force. For these images and some cases where the external force is too small, there is no slipping occurring. We run a linear regression on the remaining data points of this material to estimate p_{slip} for these data points, as shown in Fig. 5. We also add some tactile images without any contact to the dataset. For images without contact, we set $p_n = 30 \text{ mm}$, $p_{\text{slip}} = 28.5 \text{ mm} + \text{RAND}(-0.5 \text{ mm}, 0.5 \text{ mm})$, and $v_n = (p_{\text{slip}} - p_n)/3$, where $\text{RAND}(-0.5 \text{ mm}, 0.5 \text{ mm})$ generates a random distance between -0.5 mm and 0.5 mm. 3 refers to 3 seconds. This parameter controls the desired velocity of gripper width when there is no contact. A lower value might result in an excessively high desired velocity, which may not be feasible for low-level tracking. If the value is set too high, the gripper contracts too slowly when there is no contact, resulting in a significant decrease in overall system efficiency. Therefore, the choice of 3 seconds is also based on experiments tailored to our hardware specifications. During training, these images without contact allow the gripper to continue closing until it makes contact with the object.

We mount only one GelSight sensor on the follower gripper, and the other finger is with the same shape and no sensor mounted. Similarly, the two fingers on the leader gripper are all dummy fingers without mounted sensors since the leader arm is used to provide external forces during data collection. The reason for using only one GelSight on one side is because a single GelSight sensor suffices to furnish LeTac-MPC with adequate information for effective grasping. For the parallel gripper, the GelSight sensors mounted on different fingers can offer detailed information on the shape and texture of an grasped object's opposing sides, but often provide overlapping information about the information of contact force, and states and properties of the grasped object. GelSight only can perceive a very small local contact area of the object, and the idea of LeTac-MPC is using this small local area (raw tactile image) to estimate the grasping state and object property implicitly and then generate the corresponding actions. Therefore, we only mount one GelSight sensor, both for data collection and controller deployment. In addition, adding a redundant GelSight sensor would increase the computational load for real-time processing, subsequently

reducing the control frequency. In addition, we would like to highlight the following papers [16], [48]. Both of them have only one finger equipped with GelSight sensor but can accomplish complex manipulation tasks.

E. Loss Function

We propose the following loss function for our model:

$$\ell = D(\mathbf{p}_n^*, \mathbf{p}_{\text{slip}}) + \tilde{P}D(p_{n+N}^*, p_{\text{slip}}),$$

where $\mathbf{p}_n^* = [p_{n+1}^*, p_{n+2}^*, \dots, p_{n+N}^*]^T \in \mathbb{R}^N$, $\mathbf{p}_{\text{slip}} = [p_{\text{slip}}, p_{\text{slip}}, \dots, p_{\text{slip}}]^T \in \mathbb{R}^N$, and $D(\cdot, \cdot)$ is a distance measure. We use L2 distance in our implementation. The loss function measures the difference between the output predicted trajectories and the target position p_{slip} over the sequence of output positions instead of comparing a single output position to p_{slip} . This sequence-based approach promotes the convergence of the controller. Additionally, we use coefficient $\tilde{P}$ to enhance the convergence of the trained controller by increasing the terminal cost. In our implementation, we set $\tilde{P} = 3$.

Recall that in data collection, p_{slip} represents the maximum gripper width required for the gripper to hold the data collection blocks without slipping when external forces are present. Selecting p_{slip} as the target position means aiming for training a model that can output the maximum gripper width that can maintain stability of the grasped object in the hand, which corresponds to applying minimal gripping force to the grasped object.

F. Controller Implementation

The model is trained using discrete data points, while the actual tactile feedback is continuous and real-time. This difference in data formats poses significant challenges in deploying a trained model as a real-time controller with optimal performance. In our method, the proposed MPC layer combines both model-based and data-driven formulations, providing a way to fine-tune the model-based part for better control performance with real-time tactile feedback. To implement the trained model as a controller, we need to make three adjustments to the MPC layer, as shown in Fig. 6. Firstly, we increase the weight coefficients Q_v and Q_a by multiplying the scalars K_v and K_a , respectively, to increase the weights coefficients to speed up the controller's convergence. Secondly, we incorporate an additional model-based feedback signal, the marker displacement, to enhance the response speed to the tangential contact force. Finally, we add gripper motion saturation constraints in the optimization problem to ensure that control input and output are dynamically feasible. The implemented MPC controller is determined by solving the optimization problem

$$\begin{aligned} \mathbf{a}_n^* = \arg \min_{\mathbf{a}_n} & P\mathbf{f}_{n+N}^T \mathbf{Q}_f \mathbf{f}_{n+N} \\ & + PK_v Q_v (v_{n+N} - K_d d_{n+N})^2 \\ & + \sum_{k=n}^{n+N-1} \mathbf{f}_k^T \mathbf{Q}_f \mathbf{f}_k + K_v Q_v (v_k - K_d d_k)^2 + K_a Q_a a_k^2, \end{aligned} \quad (6)$$

subject to (3) and

$$\begin{bmatrix} p_{\min} \\ v_{\min} \\ a_{\min} \end{bmatrix} \leq \begin{bmatrix} p_n \\ v_n \\ a_n \end{bmatrix} \leq \begin{bmatrix} p_{\max} \\ v_{\max} \\ a_{\max} \end{bmatrix}, \quad (7)$$

where the scalar factor K_d is used to regularize d_n , which represents the length of the sum of marker displacements after applying the compensation method introduced in [49]. The scalars K_v and K_a are used to increase the weight coefficients for v and a , respectively. In our implementation, by setting $K_v = 100$ and $K_a = 2$, we are able to achieve good convergence ability and response speed for the deployed controller. Inequality (7) is the saturation constraint, where right subscripts min and max represent the lower bound and upper bound, respectively.

The reason we add constraints during the implementation of the controller instead of incorporating them during the learning phase is that an unconstrained MPC layer can better capture the features of different objects to build more applicable dynamic models. When utilizing the learned model as a controller, the application of constraints enables the entire model to better adapt to real-time tactile feedback and generate dynamically feasible motion.

IV. BASELINE METHODS

In this section, we introduce three baseline methods we use in the experiments, namely proportional derivative (PD) control, MPC, and open-loop grasping.

A. Proportional Derivative Control

Firstly, we define several control signals that can be extracted from tactile images. We denote the contact area of the grasped object with the tactile sensor as c . The contact area c can be obtained by thresholding the depth image or the difference image, as illustrated in Figs. 11 and 12. Additionally, we estimate the tangential contact force by tracking the marker displacements. We represent the length of the sum of marker displacements as d .

It should be noted that d and c here are tactile signals with clear meanings, extracted by other modules, and low-dimensional. This is distinct from $\mathbf{f}_n$ used in LeTac-MPC. $\mathbf{f}_n$ is an embedding generated by the encoder, which lacks a specific explicit meaning. During end-to-end training, we anticipate that our model would learn the physical representations of grasped objects from tactile images, which is $\mathbf{f}_n$.

We design the controller according to two objectives. Firstly, we aim to maintain a constant value c_{ref} for the contact area c reactively while grasping objects. Secondly, we




Fig. 6: Illustration of learned controller implementation. We implement the trained model as a real-time MPC controller. The proposed MPC layer combines model-based and data-driven formulations, allowing for the fine-tuning of the model-based part to achieve better control performance with real-time tactile feedback.

c_{ref}	Q_d	K_P	K_D	Δt	freq.
900	2	$\frac{1}{4 \times 10^4 \Delta t}$	$\frac{1}{3.6 \times 10^5 \Delta t}$	$\frac{1}{60}$	60 Hz

TABLE I: Parameters of the PD controller.

c_{ref}	Q_d	P	Q_c	N	K_c	Q_v, Q_a	Δt	freq.
900	2	5	36	10	36000	1	$\frac{1}{60}$	60 Hz

TABLE II: Parameters of the MPC controller.

$p_{\text{min}}/p_{\text{max}}$	$v_{\text{min}}/v_{\text{max}}$	$a_{\text{min}}/a_{\text{max}}$
0/70 mm	-15/15 mm/s	-100/100 mm/s ²

TABLE III: Constraints of LeTac-MPC and MPC.




Fig. 7: Daily objects for grasping.

p , which is a simple and valid approximation of complex gel dynamics. The linear model is a local approximation for a very short duration. In MPC, the model is not used for long-term predictions, as this would lead to significant errors. Our proposed MPC baseline method runs at a high frequency with receding horizon control, ensuring that we can continuously iterate with the local model to minimize model errors. Combining equations (1) and (8), we get the following model:

$$\begin{bmatrix} c_{n+1} \\ p_{n+1} \\ v_{n+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 & -K_c \Delta t \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} c_n \\ p_n \\ v_n \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{2} \Delta t^2 \\ \Delta t \end{bmatrix} a_n. \quad (9)$$

Defining a feedback state vector $\mathbf{y}_n = [c_n, -d_n, v_n]^T$, and the length of prediction N , we write the same control objectives with PD as following cost function:

$$J(\mathbf{y}_n, \mathbf{a}_n) = P \mathbf{e}_{n+N}^T \mathbf{Q} \mathbf{e}_{n+N} + \sum_{k=n}^{n+N-1} \mathbf{e}_k^T \mathbf{Q} \mathbf{e}_k + Q_a a_n^2, \quad (10)$$

$$\text{where } \mathbf{Q} = \begin{bmatrix} Q_c & Q_d Q_c & 0 \\ Q_d Q_c & Q_d^2 Q_c & 0 \\ 0 & 0 & Q_v \end{bmatrix}, \quad (11)$$

$$\mathbf{e}_n = \mathbf{y}_n - [c_{\text{ref}}, 0, 0]^T, \quad (12)$$

$$\mathbf{a}_n = [a_n, a_{n+1}, \dots, a_{n+N-1}]^T \in \mathbb{R}^N.$$

aim to ensure a tighter grasping in the presence of external tangential contact force to prevent the object from slipping. To achieve these goals, we define the following PD control law:

$$v_n = K_P(c_n - c_{\text{ref}} - Q_d d_n) + K_D \dot{c}_n.$$

Where K_P, K_D are gains for this controller and Q_d is a scalar factor to regularize d . The parameters we use for PD are shown in Table I.

B. Model Predictive Control

Firstly, we propose such linear assumption:

$$c_{n+1} = c_n - K_c v_n \Delta t, \quad (8)$$

where K_c is a scalar factor. Equation (8) assumes a linear relationship between the contact area c and the gripper width

In equation (10), scalar P is to amplify the terminal cost to speed up convergence. Q_c, Q_v , and Q_a are weight coefficients. Equations (11) and (12) are derived from following control objective: c_n converges to $c_{\text{ref}} + Q_d d_n$ and v_n converges to 0.

Finally, we define a MPC law by solving following opti-




Fig. 8: Illustration of daily objects grasping and transportation. Left: experiment procedures. Middle: tactile feedback sequence. Right: gripper width curve. The dynamic trajectories we use for transportation are same for every object.

	avocado	bagel	box	cracker	onion ring
LeTac-MPC	5.53 ±1.48 N	5.32 ±1.63 N	5.36 ±1.57 N	4.44 ±1.30 N	4.89 ±1.34 N
MPC	5.79±1.33 N	reach force limit	7.04±1.82 N	5.47±1.29 N	6.46±1.55 N
PD	5.96±1.40 N	reach force limit	7.85±2.00 N	6.88±1.50 N	5.83±1.37 N
open-loop	10.22±2.02 N	7.11±1.85 N	drop	9.92±1.88 N	broken
	oreo	egg	mushroom	tomato	banana
LeTac-MPC	4.67 ±1.43 N	5.43±1.42 N	4.52 ±1.51 N	5.02 ±1.70 N	5.89 ±1.68 N
MPC	5.82±1.37 N	5.28±1.34 N	5.31±1.40 N	7.42±1.66 N	9.23±1.80 N
PD	6.69±1.46 N	4.53 ±1.22 N	5.78±1.44 N	reach force limit	8.17±1.85 N
open-loop	7.78±1.89 N	13.65±2.86 N	7.02±1.75 N	7.55±1.77 N	7.35±1.83 N

TABLE IV: Results of daily objects grasping and transportation experiment. If an experiment fails, we indicate the corresponding table position with red font to highlight the failed result. If the experiment succeeds, we report the average value and standard deviation of the force during the experimental process after the grasping stabilizes. The bolded font represents the minimum grasping force for the same object. Here, the term “egg” refers to an egg with its shell, which is actually rigid. In the paper, we use “peeled boiled egg” to denote an egg without its shell.

mization problem:

$$\mathbf{a}_n^* = \arg \min_{\mathbf{a}_n} J(\mathbf{y}_n, \mathbf{a}_n), \quad (13)$$

subject to (9) and

$$\begin{bmatrix} p_{\min} \\ v_{\min} \\ a_{\min} \end{bmatrix} \leq \begin{bmatrix} p_n \\ v_n \\ a_n \end{bmatrix} \leq \begin{bmatrix} p_{\max} \\ v_{\max} \\ a_{\max} \end{bmatrix}.$$

The optimization problem (13) is a QP. Unlike the differentiable MPC layer, we do not need to do backpropagation for the optimization problem (13). Here, we only need to solve the optimization problem (13) to compute the control inputs. Therefore, we use OSQP [50], an efficient QP solver to solve the optimization problem (13) in real-time. The parameters we use for MPC are shown in Tables II and III.

C. Open-loop Grasping

For open-loop grasping, we use the force feedback of the WSG 50-110 gripper as the signal to choose the gripper

width. We select 10 N as the threshold value. At the beginning of the grasping, the gripper width decreases, and once the force feedback is greater than 10 N, the gripper stops. Then the gripper width will remain in this position during the period of performing the task.

We select 10 N as the threshold value after testing and adjusting it through grasping and dynamic transportation of the objects shown in Fig. 7 (see Section V-A for more details). 10 N is suitable for most objects, as increasing the threshold could cause damage to fragile objects, while decreasing it may result in heavy objects dropping during dynamic transportation.

It is important to note that even when we use 10 N as the static condition for the fingers in open-loop grasping, it does not imply that open-loop grasping can maintain a constant grasping force of 10 N. This variability is in part due to the inherent noise in the force feedback of the WSG gripper. Furthermore, when the force feedback reaches 10 N, the fingers are in motion, and the force exerted differs once




Fig. 9: Tactile images of daily objects grasping using LeTac-MPC.

they stop. Additionally, the final resting force varies with the shape and properties of the grasped object, as shown in Table IV.

V. EXPERIMENTS

We conduct experimental validation of proposed LeTac-MPC on a WSG 50-110 parallel gripper and a Franka Panda manipulator, as shown in Fig. 1. We run LeTac-MPC at a frequency of 25 Hz. Using the official WSG 50 ROS package¹, the maximum rate of the gripper is 30 Hz. Therefore, the gripper can track the motion generated by LeTac-MPC with reasonable saturation constraints. We perform three tasks, namely daily objects grasping and transportation (Section V-A), dynamic shaking (Section V-C), and obstacle collision (Section V-D), and compare LeTac-MPC with PD control, MPC, and open-loop grasping. To capture the tactile feedback, we mount a GelSight on one of the fingers of the WSG 50-110 gripper. A video includes these three tasks and comparisons with the baseline methods can be found in the supplemental material.

The WSG gripper we use has a built-in low-level controller, which can track position reference. Our designed

reactive grasping controllers are high-level controllers. These high-level controllers produce reference motions for the gripper fingers with reasonable saturation constraints, which are then accurately tracked by the low-level controller. In the case of LeTac-MPC and MPC, the solution to the optimization problem is acceleration a_n , but our state also includes gripper width p_n and velocity v_n . Eventually, we use p_n from the state as the reference sent to the low-level controller for tracking. Although we do not directly track v_n and a_n , since p is derived from the integration of v and a , what the low-level controller effectively tracks is the entire trajectory of the finger motion. As for the PD controller, at each moment, we calculate v_n through feedback, which, after integration, gives us the position reference p_n to be tracked by the low-level controller. Hence, the low-level controller actually tracks both p and v .

A. Daily Objects Grasping and Transportation

We evaluate our proposed LeTac-MPC method and compare it with three baseline methods for a daily objects grasping and transportation task. As shown in Fig. 7, we select 10 daily objects. These objects differ in physical properties (such as stiffness, total weight, mass distribution,

¹<https://github.com/nalt/wsg50-ros-pkg>




Fig. 10: Gripper width curves of different objects and different methods.

and coefficient of friction), sizes and shapes (ranging from larger to smaller objects), and surface textures (from smooth surfaces to different complex textures). Some objects are also fragile (e.g. onion ring, egg, and cracker). All of these present challenges to the generalization and robustness of grasping. The position of dynamic transportation is shown in Fig. 14(a). We use this trajectory for evaluating all objects and methods. The only exception is the box containing beads, which undergoes a relative displacement along the Z-axis to accommodate its larger size. Objects are handed to the gripper by someone to facilitate grasping. Among evaluating for different methods, each object will be grasped with the same configuration.

The task involves grasping different daily objects, each with different sizes, shapes, physical properties, and surface textures, and transporting them to another location. The task presents three main challenges. First, the gripper must apply the appropriate force to grasp the object without damaging it but with enough force to hold it stably. Second, during the dynamic transportation of the object, the gripper needs to reactively adjust its behavior to maintain the object stably, particularly for heavier objects and objects with varying mass distribution. Third, the task requires generalizing feasible grasping to different objects with different physical properties, shapes, sizes, and surface textures.

Based on the experimental results, our proposed LeTac-MPC overcomes these challenges and outperforms the base-

line methods. More detailed explanations of the experimental results are provided below:

1) The detailed results of the implementation of LeTac-MPC in banana and box containing beads are shown in Fig. 8. When the object was not in contact with the gripper, the gripper width decreased due to the empty tactile feedback. When the object came into contact with the gripper, the gripper width gradually converged to a position based on the tactile feedback from the GelSight. After the controller converged, the manipulator dynamically transported the object to another location, and the tactile feedback changed due to rapid changes in the states the grasped object. Correspondingly, LeTac-MPC would enable the gripper to reactively re-grasp the object to stabilize it in hand. Finally, after the object was transported to another location and became static, the gripper width converged.

Note that the contraction response of the gripper in the box experiment is faster than that of the banana. This is because the inertia of the box containing beads changes significantly during dynamic transportation, requiring a faster response to maintain its stability.

2) Fig. 9 shows the tactile images of all 10 daily objects in the experiments. Table IV shows the results and mean grasping forces of the experiments. Our proposed LeTac-MPC is able to generalize to different objects with different physical properties, sizes, shapes, and surface textures, as shown in Fig. 9. Items such as cracker (Fig. 9(d)), onion




Fig. 11: Thresholding the depth image to extract the contact area. We show a sequence of results by thresholding with different values. 0.05 is the threshold we use for PD and MPC experiments.

ring (Fig. 9(e)), and oreo (Fig. 9(f)) are smaller than the size of the gripper and also have complex textures and cluttered contours. The training dataset only includes four standardized blocks with different physical properties. This indicates a good generalizability of LeTac-MPC. In Fig. 10, we show the gripper width curves of several experiments.

3) As shown in Figs. 9(b) and 9(i), although the tactile feedback for soft objects such as bagel and tomato is subtle, LeTac-MPC can still grasp robustly. However, both PD and MPC controllers rely on the contact area as the feedback signal to compute control inputs, and therefore cannot function on soft objects that have subtle tactile feedback. This is because it is difficult to obtain a high-quality and stable contact area from tactile feedback for soft objects. There are two popular ways to extract the contact area mask of the tactile image, thresholding depth images (Fig. 11) and difference images (Fig. 12). The depth image is the 3D reconstruction of the contact surface and can be obtained by the method in [51]. The difference image is obtained by subtracting the current frame image from the first frame image, with marker removal achieved by interpolation at the

marker locations on both the current and the first frames.

We collect multiple tactile images for different objects in the experiment and present a sequence of results obtained by thresholding depth images and difference images by different values. As seen in Figs. 11 and 12, the same thresholding value cannot be generalized to different objects with different physical properties. When the value is small, we cannot obtain a reasonable contact area for deformable objects. However, when the value is large, the contact area of rigid objects becomes larger than the reasonable area and has more noises. Furthermore, for bagel, regardless of how we adjust the value, the contact area always appears poor for computing control input. This is because the bagel is very soft, resulting in very subtle tactile feedback. The raw images are from MPC experimental records. This explains the results in Table IV. For PD, since it utilizes the same feedback signals and control objectives as MPC, this can also explain the experimental results of PD.

In our implementation, we use the contact area from depth images as the feedback signal for MPC and PD. Additionally, 0.05 is the threshold that we use. As shown in Fig. 11,




Fig. 12: Thresholding the difference image to extract the contact area. To get the difference image, we need to first interpolate the raw image and remove the markers.




Fig. 13: Gripper width curves and tactile images of grasping a peeled boiled egg using LeTac-MPC.

0.05 is the optimal threshold for rigid objects with textures and without textures based on our experiments in Fig. 11. In Fig. 11, we can see that using 0.025 as the threshold introduces noise in non-contact areas for hard objects such as cracker and avocado (visible on the right side), and the contact area for cracker becomes elongated, which does not match the actual ground truth. Although tomato shows improved contact area quality at 0.025, the quality for bagel remains poor. On the other hand, a 0.05 threshold performs better for rigid objects (cracker and avocado). Therefore, we chose 0.05, which generally works for rigid objects.

The poor contact area signal leads to poor grasping behavior. As shown in Fig. 10, MPC and PD for soft object grasping continuously decreased the gripper width until it reached

the force limit of the WSG 50-110 gripper, as the large force applied by the gripper cannot create a sufficient contact area mask. However, LeTac-MPC performs exceptionally well in grasping deformable objects, even when tactile features are subtle.

Note that for tomato grasping in Fig. 10, MPC succeeded while PD did not. This does not necessarily mean that MPC outperformed PD in grasping this object. As shown in Fig. 11, the contact area in the MPC experiment is also very unsatisfactory, which results in MPC convergence, but with a large gripping force applied. Both PD and MPC calculate control inputs based on the contact area, which makes both unsuitable for grasping deformable objects. The difference in the experimental results for grasping tomato can be attributed




Fig. 14: Position trajectories for the dynamic transportation (Section V-A), dynamic shaking (Section V-C), and obstacle collision (Section V-D) experiments. The maximum velocity of the trajectory for dynamic transportation is 1.40 m/s, with a maximum acceleration of 3.82 m/s^2 . For dynamic shaking, the trajectory has a maximum velocity of 1.38 m/s and a maximum acceleration of 8.15 m/s^2 . Lastly, the trajectory for obstacle collision has a maximum velocity of 0.47 m/s and a maximum acceleration of 0.65 m/s^2 . This figure serves as a further supplement to the experiments, highlighting the dynamic nature and challenges of the trajectories involved.




Fig. 15: Gripper width curve and tactile images of dynamic shaking a box containing beads with LeTac-MPC.

to other minor factors, such as the grasping position of the tomato.

4) Table IV shows that LeTac-MPC demonstrates statistically smaller grasping forces for most objects and performs the optimal on the grasping and transportation task. For very soft objects, such as tomato and bagel, PD and MPC tend to apply excessive grasping forces. Often, these forces reach the gripper's force limit. For open-loop grasping with large grasping force, it still fails to grasp the box containing beads stably during dynamic transportation. This is because the simple open-loop policy cannot re-grasp properly when the grasped object's state changes suddenly. We can see the tactile image sequence of dropping the box in Fig. 19. Moreover, for certain objects in which the force applied during open-loop grasping is much greater than that in LeTac-MPC, we can see differences between tactile images Figs. 9 and 19. For example, by comparing Figs. 9 and 19, it is evident that the tactile image of the egg in Fig. 19 has a larger contact area and clearer features, indicating that the force applied in open-loop grasping is greater.

5) In our implementation, due to the slow convergence rate of PD and MPC, we set the gripper to contract at a constant speed of 2.5 mm/s when there is no contact with the object. When the gripper is in contact with the object, we switch

to the PD and MPC controllers. As shown in Fig. 10, the gripper width curves generated by PD and MPC both have a turning point (green circle), which represents the gripper width at the point of contact with the object. We can observe that after the turning point, PD and MPC gradually converge, but their convergence rate is slow. For LeTac-MPC, we do not set a constant contraction speed when the gripper is not in contact with an object. Instead, we apply LeTac-MPC to generate gripper motion throughout the grasping process. This is because LeTac-MPC has a faster convergence rate than PD and MPC, no matter if the gripper is in contact with the object, as shown in Fig. 10.

6) We can observe that the fluctuation after convergence of PD and MPC is smaller than that of LeTac-MPC. This is because the output of the NN has some randomness, but the output of the model-based controllers is more stable. However, as shown in Figs. 8 and 10, the randomness in the output of our proposed network model does not affect grasping performance.

B. Extreme Case: Grasping a Peeled Boiled Egg

In this section, we use LeTac-MPC to grasp a peeled boiled egg, which is extremely soft and delicate. In fact, this experiment is not successful, but shows an extreme case of




Fig. 16: Gripper width curves of dynamic shaking experiments with MPC and PD control. We know that the gripper width and grasping force are negatively correlated (smaller gripper width means tighter grasping and larger grasping force). We can see from this figure that at the beginning of the experiments, the contact areas are too small compared to the raw images. Therefore, for MPC and PD, the initial converged grasping forces are actually larger than the proper grasping force.




Fig. 17: Illustration of the obstacle collision experiment.

LeTac-MPC.

The experimental visualizations and gripper width curves are shown in Fig. 13. We can see that, during the entire grasping process, the changes in tactile images are almost imperceptible to the human eye. This is because comparing with the elastomer of GelSight, the peeled boiled egg is much more softer. So, no matter how much grasping force is applied to the peeled boiled egg, it cannot make enough deformation on the elastomer of GelSight.

When the elastomer touches the egg for the first time, there are some subtle features that show up in the tactile image. Therefore, the speed at which the gripper width decreases begins to vary, and it appears that the controller is trending toward convergence (point B in Fig. 13). However, it is important to note that the controller does not actually converge; instead, the gripper width decreases very slowly with a minimal slope (point B in Fig. 13). This is because, for grasping other daily objects in Fig. 7, when the tactile image has features, LeTac-MPC will slowly decrease the gripper width and tend toward convergence. During this slow decrease, as the grasping force increases, the features of the tactile image become more pronounced, leading LeTac-MPC to gradually converge. However, as we mentioned, since a peeled boiled egg is too soft, no matter how much force is applied to it, its tactile features do not become stronger. Therefore, the gripper width continues to decrease at a very slow rate, and LeTac-MPC never converges. This ultimately leads to the egg being damaged (points C and D in Fig. 13).

C. Dynamic Shaking

In this experiment, we have the gripper grasp a box containing beads and shake it violently to explore how LeTac-MPC and baseline methods behave to prevent the box from falling under dynamic shaking. During dynamic shaking, the beads inside also shake, resulting in unpredictable state changes in the grasped box. The position of the end-effector during dynamic shaking is shown in Fig. 14(b). The maximum acceleration can reach to 8.15 m/s^2 . This makes the task very challenging.

Fig. 15 shows the results of LeTac-MPC. We can see that during the shaking process, the box shook in the hand due to rapid changes in its acceleration and inertia. Correspondingly, the gripper re-grasped based on the tactile feedback and eventually converged. Throughout the entire process, the box did not drop even though it shook in the hand. The focus of the experiment result is on whether the box falls or not. According to Fig. 15, due to the rapid change in the state of the box, minor slips are common and we cannot control the slip distance. However, due to the rapid response of LeTac-MPC, it could quickly re-grasp the object and converge to a new grasping width to prevent it from falling.

In contrast, for open-loop control, we observe that the box easily drops because there is no reactive behavior in open-loop grasping, as shown in Fig. 19.

For MPC and PD, we observe that they have similar results, as shown in Fig. 16. We can see that, at the beginning of the experiments, the contact areas are too small compared to the raw images. Therefore, for MPC and PD, the initial convergent grasping forces are actually larger than the proper grasping force. We can also obtain this conclusion by comparing the tactile images of PD and MPC in Fig. 16 with the tactile images of LeTac-MPC in Fig. 15. In the shaking process, the gripper width first decreased due to the sudden change in the state of the grasped object, as seen in Fig. 16. Then, for both PD and MPC, the gripper width increased and converged to a new larger position value, indicating that the initial grasping forces are too large. Since the performance of model-based control methods like PD and MPC rely on the quality of tactile feedback, their generalizability is lower




Fig. 18: Gripper width curves of the obstacle collision experiment.

than that of LeTac-MPC.

We conducted 10 repetitions of the experiment for each method and calculated the success rate, as shown in Table V. We can see that the success rate of the reactive control methods (LeTac-MPC, MPC, PD) is significantly higher than that of the open-loop grasping. This is because PD and MPC tend to apply greater force due to the low-quality contact area feedback for the box, as shown in Fig. 16. Additionally, by comparing the tactile images in Figs. 15 and 16, it can be observed that while PD and MPC achieve higher success rates, LeTac-MPC maintains a relatively high success rate while applying significantly less force.

In the daily objects grasping and transportation experiment, we demonstrate that LeTac-MPC’s ability to reactively grasping to a variety of daily objects, since transportation is also a dynamic task. Additionally, this experiment also showcases both PD and MPC exhibit poor generalization across different objects. If a range of diverse objects were introduced in the dynamic shaking experiment, PD and MPC would struggle to grasp some of them effectively, particularly the soft ones. On the other hand, open-loop grasping lacks the capability for reactive grasping. In general, we anticipate that LeTac-MPC will outperform these baseline methods in such scenario. This can also be analogously applied to the obstacle collision experiment.

In Fig. 15, there is an increase to the gripper width. When shaking occurs, the state and in-hand configuration of the object may change. In such cases, the force exerted on GelSight may increase, resulting in changes in the tactile image. Consequently, the controller will increase the gripper width to mitigate this impact. This also could happen in dynamic transportation and obstacle collision experiments. However, the state of the grasped object is constantly changing. If the force exerted on GelSight decreases, the LeTac-MPC will reduce the gripper width to ensure stable grasping. The entire process is dynamic and based on tactile image feedback. Consequently, we can observe that the gripper width is sharply decreased and converges to a new position.

D. Obstacle Collision

In this section, we present the results of the experiment that tests the effect of collisions on LeTac-MPC and the baseline methods. In this experiment, we have the robot grasp a screwdriver and then collide with an obstacle, as shown in Fig. 17. This unexpected collision may occur during some robot manipulation tasks. The ideal grasping controller should be able to resist this type of external impact.

In this experiment, we chose a screwdriver as the grasped object, because it is rigid, has a regular shape, and lacks surface texture, all of which make it easily grasped by the three baseline methods (PD, MPC, and open-loop). Otherwise, as demonstrated in the daily objects grasping and transportation experiment, many objects are challenging for PD and MPC to grasp properly and robustly. Using this grasp-feasible object, we further focus on assessing the impact of obstacle collision.

As shown in Fig. 18, in the LeTac-MPC, PD, and MPC experiments, the screwdriver became loose in the hand due to the collision, resulting in the gripper rapidly contracting to re-grasp the screwdriver. Depending on the specific grasping and collision situation, reactive grasping controllers re-grasp with corresponding contracting distances and velocities. Therefore, all three methods demonstrate reactive behavior in response to this type of external impacts. On the contrary, as shown in Fig. 19, open-loop grasping cannot handle this type of external impact. When a collision occurs, the grasped object always easily drops. We conducted 10 repetitions of the experiment for each method and calculated the success rate, as shown in Table V. We can see that the success rate of the reactive control methods (LeTac-MPC, MPC, PD) is significantly higher than that of the open-loop grasping.

It is worth noting, as mentioned previously, that LeTac-MPC exhibits higher fluctuations after convergence compared to PD and MPC. However, these fluctuations are acceptable. As shown in Figs. 15 and 18, the fluctuation range is around ± 0.05 mm. As demonstrated in the supplementary video, under such minor fluctuations, changes in the tactile image are virtually indiscernible. Conversely, during dynamic shaking or when the controller reconverges after the grasped




Fig. 19: Tactile image visualizations of open-loop grasping.

object collides with an external obstacle, the changes in the tactile image become quite visible. The differences between these two indicate that within the range of ± 0.05 mm, the grasp state remains almost constant.

Additionally, the LeTac-MPC is more sensitive to changes in tactile feedback and has a faster response speed. As shown in Fig. 18, at the onset of collision, the tactile features became more pronounced, resulting in stronger feedback. At this point, the LeTac-MPC released the gripper a little bit to maintain a reasonable grasping force. However, with the change of tactile feedback, the gripper quickly contracted to stabilize the screwdriver.

VI. DISCUSSION AND FUTURE WORK

In this paper, we propose LeTac-MPC, a learning-based model predictive control for tactile-reactive grasping to address the challenges of robotic tactile-reactive grasping for objects with different physical properties and dynamic and force-interactive tasks. The proposed approach features a differentiable MPC layer to model the embeddings extracted by the NN from the tactile feedback, enabling robust reactive grasping of different daily objects. We also present a fully automated data collection pipeline and demonstrate that our trained controller can adapt to various daily objects with different physical properties, sizes, shapes, and surface

textures, although the controller is trained with 4 standardized blocks. We perform three tasks: daily objects grasping and transportation, dynamic shaking, and obstacle collision, and compare 4 methods: LeTac-MPC, MPC, PD, and open-loop grasping. Through experimental comparisons, we show that LeTac-MPC has optimal performance on dynamic and force-interactive tasks and optimal generalizability.

The fundamental idea of our LeTac-MPC lies in learning a unified representation and employing the differentiable MPC layer to utilize it. This approach enables generalization to a variety of daily objects, even when trained on a simplistic dataset. As clearly demonstrated in Fig. 9, the tactile images of each testing object show a significant divergence from the training objects. This ability to generalize to unseen objects is a challenge for other learning-based grasping methods, particularly those based on classification, to achieve. It should be noted that LeTac-MPC does not explicitly output various physical properties of the grasped object (such as stiffness). Instead, it implicitly extracts these grasp-related properties and then utilizes these implicit features through the MPC layer to generate grasping actions.

For slippery object, as long as the object is not extremely slippery, LeTac-MPC can effectively achieve reactive grasping. In our experiments, items such as a box containing beads, a tomato, and an egg, which are all slippery, are successfully

	LeTac-MPC	MPC	PD	open-loop
dynamic shaking	8/10	10/10	10/10	2/10
obstacle collision	10/10	10/10	10/10	3/10

TABLE V: Success rate of the dynamic shaking experiment and the obstacle collision experiment.

addressed by LeTac-MPC. However, this is based on the assumption that the surface of the grasped object is not extremely slippery. In such cases, the tangential force can still cause deformation of the gel surface, which can be captured by the camera in GelSight. If the surface of the object is extremely slippery and the application of sufficient grasping force does not cause enough tangential deformation of the gel surface, the algorithm lacks the necessary information to continue tightening the gripper, leading to a failure of grasping.

It should be noted that if we utilize a sensor elastomer with lower stiffness, we anticipate that LeTac-MPC, PD, and MPC will demonstrate enhanced performance in grasping soft objects. They are likely to use less force while achieving more robust grasping, attributed to the greater deformation of the elastomer upon contact with objects. However, a potential issue arises: the excessive softness of the elastomer might cause the grasped object to shake in the hand, which is undesirable for many manipulation tasks. Therefore, selecting an elastomer with the appropriate stiffness is crucial.

Our proposed LeTac-MPC has several limitations that could be potential research directions for future work.

1. First of all, our method cannot grasp very soft objects, such as various types of meat. Due to the much lower stiffness of these objects compared to the GelSight sensor elastomer, the tactile feedback is very subtle and even difficult for human eyes to recognize. To grasp such soft objects, it may be necessary to design new grippers that are more suitable for soft objects and develop sensing and control algorithms that are compatible with new grippers to solve this problem at the system level.

2. Second, our algorithm applies to grasping objects with different physical properties, but does not involve specific manipulation operations. In the future, we can explore how to perform complex manipulation tasks on objects with different physical properties based on tactile feedback.

3. In addition, our work only utilizes tactile feedback as input to the network. However, for more complex manipulation tasks, combining visual and tactile feedback can provide richer information. In such cases, the linear differentiable MPC layer may not have sufficient ability to represent task objectives. To design algorithms that can generalize to various manipulation tasks, incorporating nonlinear optimization layers with visual and tactile input could be a potential research direction.

4. Vision-based tactile sensors, like GelSight, provide high-dimensional tactile feedback. This high-dimensional feedback contains a wealth of contact information. However, processing such high-dimensional data can challenge the frame rate of the entire learning-based control loop, potentially limiting its performance in certain tasks. Tactile sensors

based on electrical signals, known for their good sensitivity and wide bandwidth, offer significant advantages. Exploring learning-based reactive grasping that combines tactile sensors based on electrical signals and vision is a promising direction.

5. Finally, since the raw image of different GelSight sensors looks slightly different, a model trained with data from one sensor might perform less effectively when transferred to another sensor. However, the imaging principle of GelSight is the same and the images from different sensors follow certain patterns and have commonalities. Therefore, training a controller or policy on one sensor and generalizing it to other GelSights presents a promising future research direction.

VII. ACKNOWLEDGEMENTS

The authors thank Yipai Du for helpful discussions about tactile sensing. The authors thank Raghava Uppuluri for setting up the initial robot system. This work was partially supported by the National Science Foundation (NSF; No. 2423068). This article solely reflects the opinions and conclusions of its authors and not of NSF.

REFERENCES

- [1] W. Yuan, S. Dong, and E. H. Adelson, "GelSight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, 2017.
- [2] S. Luo, W. Yuan, E. Adelson, A. G. Cohn, and R. Fuentes, "ViTac: Feature sharing between vision and tactile sensing for cloth texture recognition," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 2722–2727.
- [3] B. Ward-Cherrier, N. Pestell, and N. F. Lepora, "NeuroTac: A neuro-morphic optical tactile sensor applied to texture recognition," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2020, pp. 2654–2660.
- [4] M. Kaboli and G. Cheng, "Robust tactile descriptors for discriminating objects from textural properties via artificial robotic skin," *IEEE Trans. Robot.*, vol. 34, no. 4, pp. 985–1003, 2018.
- [5] B. Omarali, F. Palermo, K. Althoefer, M. Valle, and I. Farkhatdinov, "Tactile classification of object materials for virtual reality based robot teleoperation," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2022, pp. 9288–9294.
- [6] S. Luo, J. Bimbo, R. Dahiya, and H. Liu, "Robotic tactile perception of object properties: A review," *Mechatronics*, vol. 48, pp. 54–67, 2017.
- [7] J. Yin, G. M. Campbell, J. Pikul, and M. Yim, "Multimodal proximity and visuotactile sensing with a selectively transmissive soft membrane," in *Proc. IEEE Int. Conf. Soft Robotics*, 2022, pp. 802–808.
- [8] P. K. Murali, C. Wang, D. Lee, R. Dahiya, and M. Kaboli, "Deep active cross-modal visuo-tactile transfer learning for robotic object recognition," *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 9557–9564, 2022.
- [9] W. Yuan, C. Zhu, A. Owens, M. A. Srinivasan, and E. H. Adelson, "Shape-independent hardness estimation using deep learning and a GelSight tactile sensor," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2017, pp. 951–958.
- [10] S. Dong, D. Ma, E. Donlon, and A. Rodriguez, "Maintaining grasps within slipping bounds by monitoring incipient slip," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2019, pp. 3818–3824.
- [11] F. R. Hogan, J. Ballester, S. Dong, and A. Rodriguez, "Tactile Dexterity: Manipulation primitives with tactile feedback," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2020, pp. 8863–8869.
- [12] C. Lin, Z. Lin, S. Wang, and H. Xu, "DTact: A vision-based tactile sensor that measures high-resolution 3D geometry directly from darkness," *arXiv preprint arXiv:2209.13916*, 2022.
- [13] G. Zhang, Y. Du, H. Yu, and M. Y. Wang, "DelTact: A vision-based tactile sensor using a dense color pattern," *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 10778–10785, 2022.
- [14] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez, "GelSlim: A high-resolution, compact, robust, and calibrated tactile-sensing finger," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 1927–1934.

- [15] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer *et al.*, “Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3838–3845, 2020.
- [16] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, “Cable manipulation with a tactile-reactive gripper,” *Int. J. Robot. Res.*, vol. 40, no. 12-14, pp. 1385–1401, 2021.
- [17] Y. Shirai, D. K. Jha, A. U. Raghunathan, and D. Hong, “Tactile tool manipulation,” *arXiv preprint arXiv:2301.06698*, 2023.
- [18] S. Kim and A. Rodriguez, “Active extrinsic contact sensing: Application to general peg-in-hole insertion,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2022, pp. 10241–10247.
- [19] N. Sunil, S. Wang, Y. She, E. Adelson, and A. R. Garcia, “Visuotactile affordances for cloth manipulation with local control,” in *Proc. Conf. Robot Learn.*, 2022.
- [20] J. Lloyd and N. F. Lepora, “Goal-driven robotic pushing using tactile and proprioceptive feedback,” *IEEE Trans. Robot.*, vol. 38, no. 2, pp. 1201–1212, 2021.
- [21] Y. Zheng, F. F. Veiga, J. Peters, and V. J. Santos, “Autonomous learning of page flipping movements via tactile feedback,” *IEEE Trans. Robot.*, vol. 38, no. 5, pp. 2734–2749, 2022.
- [22] M. Oller, M. P. i Lisbona, D. Berenson, and N. Fazeli, “Manipulation via Membranes: High-resolution and highly deformable tactile sensing and control,” in *Proc. Conf. Robot Learn.*, 2022.
- [23] S. Tian, F. Ebert, D. Jayaraman, M. Mudigonda, C. Finn, R. Calandra, and S. Levine, “Manipulation by Feel: Touch-based control with deep predictive models,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2019, pp. 818–824.
- [24] R. Calandra, A. Owens, M. Upadhyaya, W. Yuan, J. Lin, E. H. Adelson, and S. Levine, “The Feeling of Success: Does touch sensing help predict grasp outcomes?” in *Proc. Conf. Robot Learn.*, 2017, pp. 314–323.
- [25] Z. Si, Z. Zhu, A. Agarwal, S. Anderson, and W. Yuan, “Grasp stability prediction with sim-to-real transfer from tactile sensing,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 7809–7816.
- [26] S. Kanitkar, H. Jiang, and W. Yuan, “Poselt: A visual-tactile dataset of holding poses for grasp stability analysis,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 71–78.
- [27] G. Yan, A. Schmitz, S. Funabashi, S. Somlor, T. P. Tomo, and S. Sugano, “A robotic grasping state perception framework with multi-phase tactile information and ensemble learning,” *IEEE Robot. Autom. Lett.*, 2022.
- [28] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine, “More Than a Feeling: Learning to grasp and regrasp using vision and touch,” *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3300–3307, 2018.
- [29] F. R. Hogan, M. Bauza, O. Canal, E. Donlon, and A. Rodriguez, “Tactile Regrasp: Grasp adjustments via simulated tactile transformations,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 2963–2970.
- [30] Y. Han, R. Batra, N. Boyd, T. Zhao, Y. She, S. Hutchinson, and Y. Zhao, “Learning generalizable vision-tactile robotic grasping strategy for deformable objects via transformer,” *arXiv preprint arXiv:2112.06374*, 2021.
- [31] Q. Feng, Z. Chen, J. Deng, C. Gao, J. Zhang, and A. Knoll, “Center-of-mass-based robust grasp planning for unknown objects using tactile-visual sensors,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2020, pp. 610–617.
- [32] J. W. James and N. F. Lepora, “Slip detection for grasp stabilization with a multifingered tactile robot hand,” *IEEE Trans. Robot.*, vol. 37, no. 2, pp. 506–519, 2020.
- [33] R. Kolamuri, Z. Si, Y. Zhang, A. Agarwal, and W. Yuan, “Improving grasp stability with rotation measurement from tactile sensing,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 6809–6816.
- [34] C. Wang, S. Wang, B. Romero, F. Veiga, and E. Adelson, “SwingBot: Learning physical features from in-hand tactile exploration for dynamic swing-up manipulation,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 5633–5640.
- [35] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, “Tossing-Bot: Learning to throw arbitrary objects with residual physics,” *IEEE Trans. Robot.*, vol. 36, no. 4, pp. 1307–1319, 2020.
- [36] C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, “Iterative Residual Policy: for goal-conditioned dynamic manipulation of deformable objects,” in *Proc. of Robot.: Sci. and Syst.*, 2022.
- [37] Y. Li, S. Li, V. Sitzmann, P. Agrawal, and A. Torralba, “3D neural scene representations for visuomotor control,” in *Proc. Conf. Robot Learn.*, 2022, pp. 112–123.
- [38] D. Driess, Z. Huang, Y. Li, R. Tedrake, and M. Toussaint, “Learning multi-object dynamics with compositional neural radiance fields,” in *Proc. Conf. Robot Learn.*, 2022, pp. 1755–1768.
- [39] M. Yu, K. Lv, H. Zhong, S. Song, and X. Li, “Global model learning for large deformation control of elastic deformable linear objects: An efficient and adaptive approach,” *IEEE Trans. Robot.*, 2022.
- [40] T. Bi, C. Sferazza, and R. D’Andrea, “Zero-shot sim-to-real transfer of tactile control policies for aggressive swing-up manipulation,” *IEEE Robot. Autom. Lett.*, vol. 6, no. 3, pp. 5761–5768, 2021.
- [41] W. Yuan, Y. Mo, S. Wang, and E. H. Adelson, “Active clothing material perception using tactile sensing and deep learning,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 4842–4849.
- [42] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [43] B. Amos and J. Z. Kolter, “OptNet: Differentiable optimization as a layer in neural networks,” in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 136–145.
- [44] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [45] <https://s3.amazonaws.com/mf.product.doc.images/Datasheets/Material+Datasheets/CompositesMaterialDatasheet.pdf>.
- [46] <https://static1.squarespace.com/static/5b1ecdd4372b96e84200cf1d/t/5b608e1a758d46a989c0bf93/1533054492488/XP-565+%281%29.pdf>.
- [47] <https://www.cati.com/blog/convert-durometer-to-youngs-modulus/>.
- [48] N. Sunil, S. Wang, Y. She, E. Adelson, and A. R. Garcia, “Visuotactile affordances for cloth manipulation with local control,” in *Proc. Conference on Robot Learning*, 2023, pp. 1596–1606.
- [49] D. Ma, E. Donlon, S. Dong, and A. Rodriguez, “Dense tactile force estimation using gelslim and inverse FEM,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2019, pp. 5418–5424.
- [50] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, “OSQP: an operator splitting solver for quadratic programs,” *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020. [Online]. Available: <https://doi.org/10.1007/s12532-020-00179-2>
- [51] S. Wang, Y. She, B. Romero, and E. Adelson, “GelSight Wedge: Measuring high-resolution 3D contact geometry with a compact robot finger,” in *Proc. IEEE Int. Conf. Robot. Autom.*, 2021, pp. 6468–6475.



Zhengtong Xu received his Bachelor’s degree in mechanical engineering from Huazhong University of Science and Technology, China. He is currently pursuing his Ph.D. at Purdue University. His research focus is on robot learning.



Yu She is an assistant professor at Purdue University School of Industrial Engineering. Prior to that, he was a postdoctoral researcher in the Computer Science and Artificial Intelligence Laboratory at MIT from 2018 to 2021. He earned his Ph.D. degree in the Department of Mechanical Engineering at the Ohio State University in 2018. His research, at the intersection of mechanical design, sensory perception, and dynamic control, explores human-safe collaborative robots, soft robotics, and robotic manipulation.