

# Learning Generalizable Vision-Tactile Robotic Grasping Strategy for Deformable Objects via Transformer

Yunhai Han<sup>1</sup>, Rahul Batra<sup>1</sup>, Nathan Boyd<sup>1</sup>, Tuo Zhao<sup>2</sup>, Yu She<sup>3</sup>, Seth Hutchinson<sup>1</sup>, and Ye Zhao<sup>1,\*</sup>

**Abstract**—Reliable robotic grasping, especially with deformable objects such as fruits, remains a challenging task due to underactuated contact interactions with a gripper, unknown object dynamics and geometries. In this study, we propose a Transformer-based robotic grasping framework for rigid grippers that leverage tactile and visual information for safe object grasping. Specifically, the Transformer models learn physical feature embeddings with sensor feedback through performing two pre-defined explorative actions (pinching and sliding) and predict a grasping outcome through a multilayer perceptron (MLP) with a given grasping strength. Using these predictions, the gripper predicts a safe grasping strength via inference. Compared with convolutional recurrent networks, the Transformer models can capture the long-term dependencies across the image sequences and process spatial-temporal features simultaneously. We first benchmark the Transformer models on a public dataset for slip detection. Following that, we show that the Transformer models outperform a CNN+LSTM model in terms of grasping accuracy and computational efficiency. We also collect a new fruit grasping dataset and conduct online grasping experiments using the proposed framework for both seen and unseen fruits. Our codes and dataset are public on GitHub<sup>1</sup>.

**Index Terms**—Deep Learning, Visual and Tactile Sensing, Perception for Grasping and Manipulation.

## I. INTRODUCTION

ROBOT manipulation has been widely used in industries for decades, but mostly for repetitive tasks in structured environment where there is little uncertainty or contact deformation in manipulated objects. For the tasks where object contact parameters are prone to vary, such as fruit grasping, they are still challenging for robotic systems [1]. Loose grips with small grasping forces can cause objects to slip, while large grasping forces can cause damage. Additionally, object contact geometry and frictional properties may also affect the optimal grasping forces for safe grasping. To learn general-purpose grasping skills, robots need to leverage with dense notions of contact information from in-hand interactions.

To model the dynamic interactions between the object and its environment, vision-based sensing frameworks have been studied based on a sequence of visual observations obtained by external cameras [2], [3]. However, these methods are not sensitive to the dense local deformation near contact




Fig. 1: Demonstration of our robotic grasping experimental set-up. The robot gripper safely grasp the fruits on the table and sorts them into the target bins via the learned framework. Robot setup: a KUKA LBR iiwa robot is equipped with a gripper (red box), of which both fingers are equipped with a GelSight sensor (blue circle). A Realsense D435 is mounted above the gripper (green ellipsoid).

regions, which could lead to errors between the perceived and actual states of a grasp. To address this issue, tactile sensing has gained increasing popularity recently [4]. Among various tactile sensors, the ones with internal cameras, such as GelSight sensor [5], have the capability of capturing high resolution image data regarding local contact geometry. Other tactile designs [6], [7] have also demonstrated a variety of manipulation tasks with similar methods. Compared to force sensors, tactile sensors can capture an object's deformation during contact. Moreover, tactile data can be readily integrated by modern learning methods for classification and task-oriented control policy learning [8]. In [9], [10], they demonstrated that grasping performance can be significantly improved by incorporating visual and tactile sensing.

In this paper, we employ two state-of-the-art Transformer models – TimeSformer [11] and ViViT [12] – to determine safe grasping forces from the visual and tactile image sequences collected during pre-designed explorative actions (e.g., pinching and sliding). The idea of designing task-oriented explorative actions is inspired by [8] and the motivations of introducing the Transformer models are: 1) compared with recurrent networks, such as LSTM, they do not suffer from the forgetting issue, 2) compared with convolutional networks used for extracting local features, they have larger receptive fields that are helpful to understand the global context, and 3) compared with CNN+LSTM models for processing image sequences, they can extract the spatial-temporal features simultaneously. While for CNN + LSTM models, the per-frame

<sup>1</sup>Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30332, USA (email: {yhan389, gtg693m, nboyd31, seth, yzhao301}@gatech.edu).

<sup>2</sup>School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA (email: tourzhao@gatech.edu).

<sup>3</sup>School of Industrial Engineering, Purdue University, 610 Purdue Mall, West Lafayette, IN 47907, USA (email: shey@purdue.edu).

\* Corresponding author

Special credit to GelSight hardware support from Edward Adelson's lab at MIT.

<sup>1</sup><https://github.com/GTLIDAR/DeformableObjectsGrasping.git>

spatial features are always encoded (CNN) prior to the temporal decoding (LSTM). Thus, the Transformer models are more adaptable to complex tasks. In our framework, the Transformer models learn low-dimensional embeddings in a supervised fashion for each sensor modality and then output a fused physical feature embedding. Firstly, we take this embedding as input and combine it with a given grasping force threshold to predict the final grasping outcomes through a multilayer perceptron (MLP). The grasping outcomes are categorized into three labels: safe grasping, slippery, and potential damage. A force threshold for safe grasping is then searched for using the learned predictor during online deployments. Secondly, the fused physical feature embedding is used to classify grasped fruit types through a different MLP layer in order to place them into separate bins automatically.

To validate the grasping framework, we perform grasping experiments on various deformable fruits for data collection. We train the models using both camera and GelSight inputs and test its performance via grasping outcome classifications on unseen fruits and online grasping success rate for both seen and unseen fruits. In addition, we benchmark the Transformer models against a CNN+LSTM model on a public dataset for slip detection [9]. Both Transformer methods (TimeSformer and ViViT) outperformed the CNN+LSTM model by 3.1% and 2.0% in detecting slip, and are much computationally efficient, making them more suitable for online tasks.

The contributions of our work are summarized as follows:

- We propose a Transformer-based grasping framework for fruit grasping and demonstrate the superior efficacy and efficiency of the Transformer models against a CNN+LSTM baseline model.
- We design a learning-based control framework that incorporates safe grasping force estimation using tactile & visual information obtained via two explorative actions: pinching and sliding, which do not require any prior knowledge of physical contact or geometrical models. Besides, the control parameter is directly formulated as the depth value read from the tactile feedback without any aid of external force-torque sensors.
- We experimentally evaluate the proposed grasping framework on a diverse set of fruits and achieve an end-to-end demonstration of fruit grasping. Besides, by performing the attention analysis, we show that the trained Transformer models take advantage of the attention mechanism to: i) incorporate more contact area information for the grasping task, such as local contact region in tactile images and fruit surface near gripper's fingertips in visual images; and to ii) capture long-term dependencies between initial and final grasping status.

## II. RELATED WORK

### A. Robotic Grasping

Robotic grasping has been a widely explored topic using numerous gripper designs and sensor modalities [13]. Recently, motivated by human's intense dependence of tactile feedback for grasping process, tactile sensors have thus begun to play an important role in robotic grasping [14]. In [15],

the authors used deep-learning methods to obtain a grasping policy for rigid grippers. However, the grasping success rate on deformable objects was not ideal since they only adjusted the grasping position but fixed the grasping force. The work of [16] estimated the optimal grasping force empirically but assumed the object weights were known. In [17] [18], the gripper's opening were controlled to stabilize the grasped objects under external disturbances by detecting the slip occurrences. Similarly, the work in [19] made use of tactile sensing to stabilize the grasped object by controlling the grasping force. However, all of these studies assumed that the objects were already steadily grasped in hand. In this work, we aim at estimating safe grasping force for deformable objects through a learning framework.

### B. Vision-tactile Sensor Fusion

We can improve the manipulation performance by fusing the information obtained from visual and tactile sensors. In [10], they proposed a multi-modal sensing framework for grasping outcome prediction. Their subsequent work in [15] investigated a learned regrasp policy based on visuo-tactile data after executing an initial grasp. Their results indicated that incorporating tactile readings substantially improves the grasping performance. However, the manipulated objects used in their experiments are primarily rigid objects, which do not require accurate force control. In other works [9], [20], [21], they used CNN + LSTM models to classify the slip occurrence, to recognize the object instance, and to perceive the physical properties of objects. Nonetheless, these methods can only be used for classification tasks and not applicable to learn control policy for safe manipulation.

### C. Transformers for Robotics

Transformer models were originally proposed for natural language processing (NLP) [22] and computer vision (CV) [23][11][12]. Recently, Transformers have drawn increasing attention in robotics. The authors of [24] proposed a Transformer framework for tabletop tasks, which encode language goals and RGB-D voxel observations and output discretized 6-DoF actions. In [25], they explored the use of Transformers to predict robot action commands for accurate object throwing. The study of [26] addressed quadrupedal locomotion tasks using Reinforcement Learning (RL) with a Transformer-based model. All of these works showed significant improvements over baseline methods on task performance and training efficiency. However, to the best of our knowledge, no existing study has ever explored the use of Transformers for robotic grasping using tactile and visual images.

## III. METHODS

In this section, we describe the details of the grasping framework and each Transformer model. In order to give robots the ability to estimate the safe grasping force, we first let the robot obtain physical information about the target objects (fruits in this work) by performing two explorative actions, **pinching** and **sliding**, on the objects. To avoid any potential damage, these actions have minimum interaction with the objects. To monitor the interactions and record the

data, the robot is equipped with two different sensors. Next, a force threshold for safe grasping will be searched for via inference using the obtained physical information and adopted for execution. In the following, we first describe the sensors in Section III-A, and then we discuss the Transformer models in Section III-B, and finally we propose the grasping framework in Section III-C and Section III-D.

#### A. Sensing Modalities

1) *Tactile*: The GelSight [5] sensor provides the robot with dense visual information (high resolution image) about the contact region between the objects and the robot fingertips. For this purpose, the contact surface of the sensor is covered with a soft elastomer such that the sensor can measure the object's compliance by observing the elastomer's vertical and lateral deformation. In our experiments, the gripper has two GelSight sensors installed on the fingertips, but we only use one to demonstrate a minimum system setup.

2) *Vision*: A RealSense D435 camera is used in this work, and we only consider the RGB datastream. The camera is wrist mounted at an angle of 15 degrees such that the image is centered on grasped objects (see Fig. 1 for the setup).

#### B. Transformer Model

We apply the Transformer models for two robotic manipulation tasks: slip detection and safe grasping force estimation. For slip detection, we replace the CNN + LSTM model used in [9] with the Transformer models but keep the last Fully Connected (FC) layer with two outputs (i.e., a stable grasp or slip) as the final classification results. For safe grasping force estimation, the outputs from Transformer models are used as inputs to the subsequent models in the grasping framework that will be thoroughly described in Section III-C & Section III-D.

Two lightweight Transformer models are explored for robotic tasks in this work: **TimeSformer** [11], **ViViT** [12]. For each model, the self-attention mechanism is similar, while the difference lies in the factorizing strategy for the spatial-temporal attention.

1) *Self-Attention Mechanism*: The self-attention mechanism allows all the inputs to be interacted with each other and identify the ones that should be paid more attention to, which brings the main advantages over CNN+LSTM models. Specifically, this mechanism can be described as mapping a query ( $Q$ ) and a key ( $K$ )-value ( $V$ ) pair to the outputs.

First, for a single self-attention block (or a single head), the query, key, and value vectors can be computed by projecting the same input matrix  $X \in \mathbb{R}^{n \times d_x}$  (each row of  $X$  corresponds to an input vector with size  $d_x$ ) to  $Q, K, V$  as follows:

$$Q = XW^Q, K = XW^K, V = XW^V \quad (1)$$

Here,  $W^Q \in \mathbb{R}^{d_x \times d_k}$ ,  $W^K \in \mathbb{R}^{d_x \times d_k}$  and  $W^V \in \mathbb{R}^{d_x \times d_v}$  are learnable matrices and  $d_k = d_v = d_x$ .

Next, the outputs are obtained through Eqn. 2, which are the weighted sums of the value vectors and the weight assigned to each column in  $V$  is a compatibility function of  $Q$  with the corresponding  $K$  at the same vector index. The dot-product is scaled by  $\sqrt{d_k}$ , as suggested in [22].

$$\text{Attention}(Q, K, V) = \text{softmax} \left( \frac{QK^T}{\sqrt{d_k}} \right) V \quad (2)$$

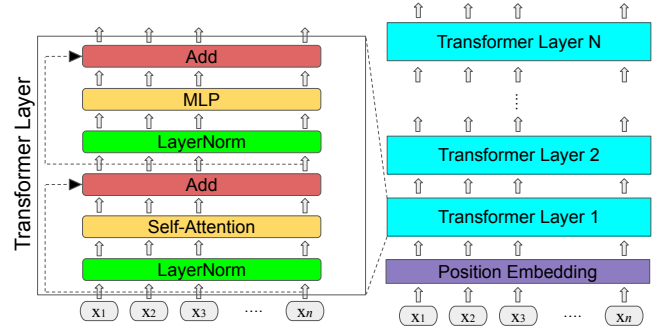


Fig. 2: An illustration of the Transformer model structure. The left figure shows one Transformer layer and the right figure shows the encoding structure.  $x_1, x_2, x_3, \dots, x_n$  are the input vectors which are first linearly embedded and then added with the position embeddings before Transformer layer 1.

$$\text{SingleHead}(X) = \text{Attention}(Q, K, V)W^O \quad (3)$$

Then, as shown in Eqn. 3, another learnable matrix  $W^O \in \mathbb{R}^{d_v \times d_x}$  projects the intermediate results to the new matrix with the same dimension as  $X$ .

In practice, instead of performing a SingleHead function, it is beneficial to project  $Q, K, V$  matrices  $h$  times with different sets of weights  $\{W_i^Q, W_i^K, W_i^V\}_{i=1}^h$  because it allows the model attend to information from different combinations of input space representations [22]. As a result, the MultiHead strategy is always employed for the Transformer models.

In addition to the self-attention layer, there is a fully-connected MLP layer that is applied to each vector position separately and identically. It has two linear transformations and a GeLU activation function in between.

A Transformer layer contains a self-attention layer and a MLP layer. To stack the Transformer layers for a deeper encoding structure, the MLP layer does not change the vector size. Also, before and after each layer, there is a LayerNorm and a residual connection, respectively. One Transformer layer is shown in Fig. 2 (left-side subfigure), where the outputs of the current layer will be the inputs for the next (right-side subfigure). Before the first Transformer layer, all the input vectors will be linearly embedded and then added with position embeddings, the elements of which represent the positions of each vector, to retain the useful sequence knowledge [22].

2) *Factorization of Spatial-Temporal Attention*: For image-based tasks, to generate the input vectors from the raw image(s), in [23], they split an image into fixed-size patches and embed each of them via linear transformation. Our framework for safe grasping handles image sequences instead of single images and must consider the temporal dimension within each self-attention layer. To accomplish this, we incorporate spatial-temporal factorization using TimeSformer [11] and ViViT [12].

**TimeSformer**: In this model, spatial-temporal dimensions are processed sequentially: within each self-attention layer, the attention is first applied on the temporal dimension of the inputs at the same spatial position, followed by the spatial dimension among all inputs from the same temporal position. There are also residual connections between each operation.

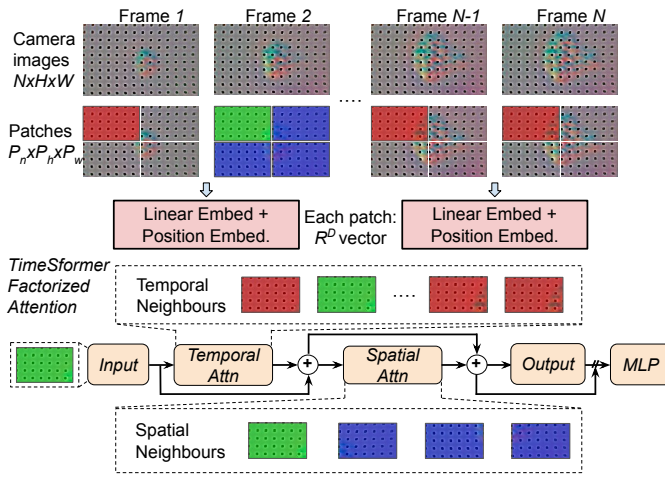


Fig. 3: Visualization of space-time attention approach for TimeS-former. The top three rows show an input GelSight image sequence, the generated  $4N$  image patches and the patch embeddings. We denote one image patch in green and its spatial-temporal neighbours in blue and red, respectively. Within each self-attention layer, for each image patch, the attentions across temporal neighbours and spatial neighbours are sequentially processed and the output will be the input of the MLP layer in Transformer model after LayerNorm.

This approach is visualized in Fig. 3. In our work, the input image sequence is denoted as  $X_I \in \mathbb{R}^{N \times H \times W}$ , where  $N, H, W$  are the number of images, image height pixels, and image width pixels, respectively. We first extract the patches  $X_P \in \mathbb{R}^{P_n \times P_n \times P_w}$ , where  $(P_n, P_w)$  is the resolution of each patch and  $P_n = \frac{NH}{P_n}$ . Next, these patches are flattened and then linearly embedded to vectors of size  $D$  with a positional embedding being added to each of them. We further add a CLS (classifier) token to the sequence of embedded vectors, which is designed to extract task-level representations [27] by attending to all the other vectors. Therefore, we can obtain the input matrix as  $X \in \mathbb{R}^{(P_n+1) \times D}$  and it then feeds into a series of Transformer layers. Finally, the output of the CLS token (size  $D$ ) from the last Transformer layer is used for different tasks. For slip detection (Section. IV-A), it is passed through a MLP layer to classify whether or not a slip occurs.

**ViViT:** Our implementation of ViViT is similar to TimeS-former, except for the following differences: First, both dimensions are processed in parallel. Specifically, half of the heads attend to the spatial dimension and the other half to temporal dimension (factorised dot-product attention). We then combine each output by concatenation and add a linear transformation to halve the size. Second, there is no CLS tokens added to the embedded input vectors because of the ambiguities when dot-producting the temporal and spatial attentions. Instead, we take the average of all patch outputs from the last Transformer layer and pass it (size  $D$ ) to the MLP layer for slip detection.

### C. Grasping Framework for Safe Force Estimation

The main goal of our grasping framework is to predict the grasping outcome given a grasping force threshold and to estimate the force threshold for safe grasping via inference.

1) *Grasping Outcome Prediction:* As shown in Fig. 4, this framework is composed of five main components: Control Pa-

rameter (Force Threshold), Transformer, Sensor Fusion model, Action Fusion model, and Prediction model.

**Force Threshold:** GelSight is a vision-based tactile sensor, which lacks the capability of estimating the grasping force (contact normal force) directly. To address this issue, the authors in [6] showed that the contact normal force can be estimated from the depth value (unit: pixel) with accurate gel calibration. On the other side, the work in [28] directly used the mean value of the marker displacement provided by GelSight images to estimate the resultant frictional force. Inspired by them, we employ the maximum depth value as the approximation of grasping forces. If the maximum depth value feedback is larger than the selected threshold for three continuous steps when running the framework, the gripper will begin to grasp the fruit. Also, the force threshold will be sent into the Prediction Model. Note that the unit of force threshold is pixel, which will be omitted in Sec. IV for readability.

**Transformer:** For each explorative action, the image sequence from a sensor modality outputs a vector of size  $D$  via the Transformer models, as thoroughly described in Sec. III-B. In Fig. 4, since we pre-define two explorative actions and there are two sensor modalities, we have 4 vectors:  $v_{\text{visual}}^{\text{pinch}}, v_{\text{tactile}}^{\text{pinch}}, v_{\text{visual}}^{\text{slide}}, v_{\text{tactile}}^{\text{slide}}$ .

**Action Fusion model & Sensor Fusion model:** We concatenate each two vectors obtained from the same exploration action and achieve:  $v^{\text{pinch}} = [v_{\text{visual}}^{\text{pinch}}, v_{\text{tactile}}^{\text{pinch}}]$  and  $v^{\text{slide}} = [v_{\text{visual}}^{\text{slide}}, v_{\text{tactile}}^{\text{slide}}]$ . We then fuse them as a vector of size  $4 \times D$  and project it to a low-dimensional space, which is considered as a fused physical feature embedding.

2) *Safe Force Threshold Estimation:* We aim to identify the control parameter, i.e., the safe grasping force threshold. As shown in Fig. 4, the Prediction Model takes the low-dimensional physical embedding obtained from performing two explorative actions and a force threshold candidate as inputs and outputs the upcoming grasping outcome via learnable neural network layers. Next, we can either uniformly or randomly sample the thresholds and feed each of them into the prediction model and select the one that predicts a safe grasping. When there are multiple viable choices, we select the average value.

### D. Grasping Framework for Fruit Classification

The second goal of our grasping framework is to classify the grasped fruit type for fruit picking operations. Specifically, we use another MLP network as a Fruit Classification model which takes the fused physical feature embedding of Sensor Fusion model as input and outputs the labels of the grasped fruit type (**Fruit Type** Block in Fig. 4). During training, the weights of the Transformer models that generate the embedding are frozen and only the MLP network is trained.

## IV. EXPERIMENTS

In this section, we present our experiments using the Transformer models. First, we benchmark the Transformer models against a CNN+LSTM model on a public dataset for slip detection with different sensor modalities. We then examine our framework for grasping deformable fruits. The robot setup is shown in Fig. 1.

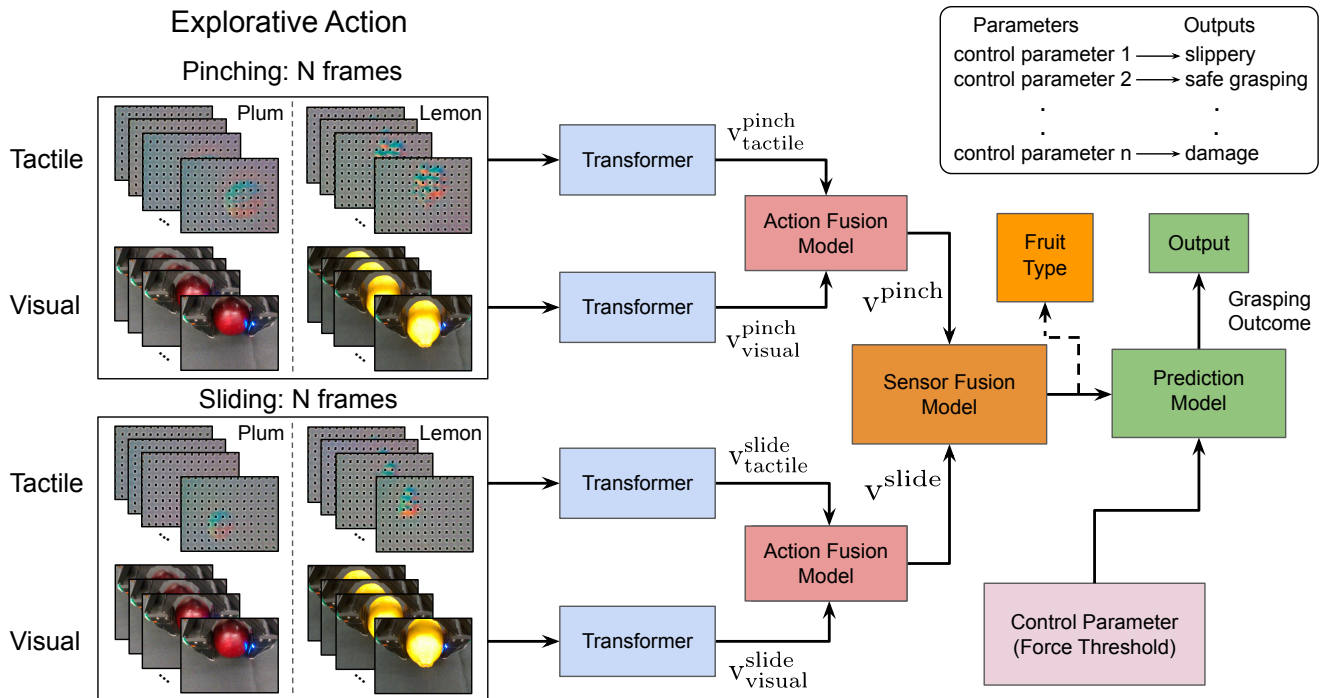


Fig. 4: **Overview of Grasping Framework.** The robot first performs two explorative actions: 1). pinching the fruit, 2). sliding along the fruit surface in the optical axis. Many fruit examples can be found on our GitHub page. Each image sequence is processed by an individual Transform network into a vector of size  $D$ . The fusion models concatenate these vectors and project it into a low-dimensional fused physical feature embedding. This embedding is further processed by a Fruit Classification model to classify the grasped fruit type. Besides, the prediction model takes the same embedding and control parameter (force threshold) as inputs and predicts the final grasping outcome. Through inference, a set of control parameters is first generated and then the parameter with the safe grasping outcome is selected to perform online grasping. This procedure is shown in the top-right black box. If there are multiple viable choices, we select the average value.

### A. Transformers for Slip Detection

1) *Experiment Setup:* We first conduct experiments for a slip detection task. The dataset released by [9] is used and can be directly downloaded and can be directly downloaded online<sup>2</sup>. We then compare the detection accuracy of the Transformer models with their method. For implementation, the whole dataset is split into training, validation, and test sets, where the test data uses unseen objects in the training data. Since the size of dataset is relatively small, we randomly split the whole dataset five times and train the model on each of them to mitigate the effect of overfitting. The final detection accuracy on the test set is averaged. For each model, we analyze the performance with three different data source inputs (vision-only, tactile-only, and vision & tactile). For the CNN+LSTM model, ResNet18 is chosen as the CNN architecture over other options, such as VGG or Inception [9], due to its advantage of fewer parameters. As a result, the ResNet18 architecture can be initialized randomly without the need of loading a pre-trained model.

A sequence of 14 continuous frames are used as input for each sensor data. During training, we use cross-entropy (two categories) as the loss function and apply an Adam optimizer [29]. For both Transformer models, the input embedding size ( $D$ ), number of Transformer layers, and number of heads are set to be 256, 8, 16, respectively. The experiment results and execution time are shown in Table I.

Accuracy \ Model	CNN + LSTM ResNet18	TimeSformer	ViViT
Modality			
vision-only	71.7% (0.4%)	78.7% (0.7%)	78.9% (1.2%)
tactile-only	80.6% (0.8%)	81.0% (0.5%)	81.8% (0.5%)
vision & tactile	81.9% (0.3%)	<b>85.0%</b> (0.4%)	83.9% (0.3%)
Execution time of feed-forward test	9.61s	2.46s	2.43s

TABLE I: Experimental results on slip detection dataset [9]. TimeSformer and ViViT outperform the CNN + LSTM method by 3.1% and 2.0%, respectively. Recorded values are the average across 5 dataset splits and their variances in parenthesis.

2) *Experiment Analysis:* From Table I, we can see that the Transformer models can provide more accurate classification results. Also, when tested on the same dataset split as in [9], the Transformer models can achieve better results (92.3% for TimeSformer and 90.0% for ViViT) than reported in [9] (88.0%) using both sensor inputs. One potential reason for the efficacy of Transformer models is that in this application, the final grasping outcomes may be inferred partially from the initial grasping status and Transformer models have the capacity of capturing these long-term temporal dependencies more effectively compared with recurrent networks [30]. Another potential reason is related to the architecture of the Transformer models: since each Transformer layer is stacked in a sequence, spatial and temporal information can be extracted simultaneously via self-attention mechanism, which does not hold for the CNN + LSTM models.

In addition, it takes significantly less time for feed-forward computation of the trained networks during the robotic deploy-

<sup>2</sup>[https://drive.google.com/file/d/1\\_iTPx18TENznXVh-82126kW9qXrNsV-/view?usp=share\\_link](https://drive.google.com/file/d/1_iTPx18TENznXVh-82126kW9qXrNsV-/view?usp=share_link)




Fig. 5: Top row: the fruits used during experiments. From the left to the right, they are: plums, apples, lemons, tomatoes, oranges and kiwifruits. Bottom row: the Gelsight images collected at the final frame during pinching for each fruit grasping. It can be seen that the fruit deformation sensed by Gelsight varies as they share different hardness and surface texture.

ment. As shown in Table I, using both vision & tactile inputs from the same test dataset and selecting the same batch size, the execution time of both TimeSformer and ViViT models on the same machine (NVIDIA GeForce RTX 2070) is 2.46 s (25.6%) and 2.43 s (25.3%), compared to the CNN + LSTM model (ResNet18: 9.61 s, VGG16: 21.39 s). Therefore, it can be concluded that Transformer enables the robots to make decisions within a much shorter time.

For the CNN+LSTM model, tactile-only significantly outperforms vision-only, as also reported in [9]. The Transformer models perform similarly for each single sensor case while also show better performance for the multi-sensor case. This indicates that multi-sensor input provides better cues for the slip detection.

### B. Transformers for Safe Fruit Grasping

1) *Experiment Setup*: We collect our own dataset<sup>3</sup> on fruit grasping involving six different types of fruits: plums, oranges, lemons, tomatoes, apples, and kiwifruits, as shown in Fig. 5 (top row). We perform the fruit grasping with various grasping force thresholds (discussed in SecIII-C1) on plum, orange, lemon, tomato, apple and kiwifruit for 782 times in total to train the models. Each fruit is clearly visible with a black backdrop relative to the camera frame. On the bottom row in Fig. 5, we show that the fruit deformation varies during pinching as they differ in hardness and surface texture. The data is collected by the RealSense camera and GelSight at 30 Hz and with  $640 \times 480$ ,  $200 \times 150$  resolutions, respectively. The visual images are then resized to  $160 \times 120$  resolution for computational efficiency. For both pinching and sliding actions, we use the *first frame of every three continuous frames* for a total of 8 frames (frame index: 1, 4, 7, 10, 13, 16, 19, 22). For both Transformer models, the patch sizes are set as (20,15) for tactile data and (16,12) for visual data. The input embedding size ( $D$ ), number of Transformer layers, and number of heads are set to be 256, 16, 8, respectively.

2) *Model Evaluation*: First, we ask can our Transformer models outperform the CNN + LSTM models in terms of grasping outcome prediction for unseen fruits? Thus, we use **cross-validation technique**, partitioning one type of fruit grasping data as testing set, to compare the model performance. After training, the Transformer models achieve 80.2% and 76.0% accuracy of grasping outcome prediction on the

Success Rate \ Model	Model		
	CNN + LSTM ResNet18	TimeSformer	ViViT
Plum	66%	88%	86%
Orange	64%	86%	90%
Lemon	60%	90%	90%
Tomato	74%	86%	86%
Apple	68%	92%	92%
Kiwifruit (unseen)	52%	74%	80%
Computation time for one sample	0.52 s	0.39 s	0.29 s

TABLE II: Experimental results of success rate for online fruit grasping (50 trials for each fruit) and the average computation time.




Fig. 6: The times each force threshold candidate is selected for safe grasping. The positional order of the color bars at each sample is shown in the dashed block.

test dataset (kiwifruit) for TimeSformer and ViViT, respectively. For CNN + LSTM model with Resnet18 as the CNN architecture, it achieves 75.0% accuracy on the test dataset.

3) *Online Fruit Grasping Evaluation*: Then, the trained frameworks are deployed on a 7-DOF KUKA LBR iiwa robot manipulator to estimate the safe grasping force via inference for both seen and unseen fruits. During inference, we sample the force thresholds as integers between 4 and 16, and for each sample, we adopt the same fused physical embedding obtained from performing two pre-defined explorative actions. The robot then grasps each fruit 50 times. Table. II shows the success rate and the average on-board computation time for one sample. It can be seen from Table. II that the Transformer models outperform the CNN+LSTM model significantly, for both seen and unseen fruit grasping.

It is noteworthy that in spite of the variations in grasping position and fruit status caused by manual fruit reloading and slight fruit spoilage, the framework is still able to select the safe grasping threshold for each grasp. This shows that the framework demonstrates some level of generalizability under the uncertainty of the local contact geometry and fruit ripeness. Besides, Fig. 6 shows the times each force threshold candidate is selected for the successful fruit grasping when using ViViT. Their values are proportional to the grasping force that should be exerted on the fruit. It is observed that the selected force thresholds for safe grasping of each fruit distribute over a finite range. Take orange as an example, force threshold 4.5 is selected 11 times for safe grasping of softer oranges and 6.0 is also selected 5 times for harder oranges. This force threshold variation potentially indicates our framework's adaptation to

<sup>3</sup>[https://drive.google.com/file/d/144cLi-LkPZVHD\\_JHfJSC8mYAk6gISCNI/view?usp=share\\_link](https://drive.google.com/file/d/144cLi-LkPZVHD_JHfJSC8mYAk6gISCNI/view?usp=share_link)




Fig. 7: Visualization of temporal attention from selected image patches at the final frame to their temporally preceding neighbours during a lemon grasping. We only show the results of four frames here. From top to bottom rows, the images are collected from pinching (tactile), sliding (tactile), pinching (visual) and sliding (visual), respectively. The image at each frame is split into  $10 \times 10$  patches, among which 24 and 6 patches are selected (denoted within the dashed ellipsoid in final frame) to present the temporal attention flows for tactile and visual images, respectively. The brighter the patch color is, the more attention is paid to the patch from its temporal neighbour at final frame.




Fig. 8: The normalized temporal attention weights to all selected image patches at each frame from their corresponding temporal neighbours at final frame.

the fruit’s inherent variability, which is infeasible by hard coding a fixed force threshold, even for the same fruit.

4) *Attention Analysis*: The most important component of our framework is the Transformer models learned entirely from data. Therefore, we now examine what pattern has our models learned qualitatively. Take TimeSformer as an example, we use the Attention Rollout method [31] to visualize the learned temporal attentions across vision & tactile image sequences on several selected image patches, as shown in Fig. 7. It can be seen that the image patches at the final frame do not only attend to themselves, but also their temporal neighbours at preceding frames (red color brightness denotes the attention weights). In addition, Fig. 8 shows the normalized temporal attention weights of all selected image patches. An intriguing observation is that the image patches at first two frames, when the gripper initially touch the objects, share larger attention weights compared with succeeding interme-




Fig. 9: Visualization of spatial attention from output token to the input image space at Frame 22 during fruit grasping. From top to bottom, the images are collected from pinching (tactile), sliding (tactile), pinching (visual), and sliding (visual), respectively. The image brightness corresponds to spatial attention weights. It is clear that the model mostly attends to the **local contact region** on tactile images and attends to the **fruit surface near gripper’s fingertips** on visual images. It should be noted that kiwi is unseen during training.

diated frames. Our conjecture of this observation is due to the fact that as the initial and ending contact information is more inferable to the physical status of the manipulated objects, as well as to the grasping outcome. On the contrary, the gradient flow in recurrent networks, even for LSTM architecture, can gradually lose information of the previous inputs, especially of the first few inputs, resulting in the difficulty of capturing long-term temporal dependencies [32]. However, Transformer is able to mitigate this problem as demonstrated.

Furthermore, the spatial attentions at the final frame for apple, plum, (seen during training) and kiwifruit (unseen during training) grasping are shown in Fig. 9. For tactile images, the TimeSformer model mostly attends to the local contact region, and for visual images, it attends to the fruit surface near gripper’s fingertips. Therefore, the Transformer models can incorporate more contact information for the grasping task.

We highlight that the interpretability of attention mechanisms may provide with an alternative way of analyzing how deep learning methods understand the object’s physical deformation properties captured by tactile and visual sensors during robotic contact-rich tasks.

5) *Fruit Picking Experiment*: Lastly, we test the Fruit Classification model on the seen fruits during online deployments (video can be found here <sup>4</sup>), which enables the KUKA LBR iiwa robot to place each fruit into separate bins using its built-in motion controller after successful grasp. For this, we pre-define five different waypoints for each fruit and when the robot reaches the desired waypoint, it would drop off the

<sup>4</sup><https://www.youtube.com/watch?v=W7o8DsTivTk>




Fig. 10: Snapshots of fruit picking operation. (a): Grasp; (b): Move; (c): Place. Other fruit experiments are shown in the attached video.

grasped fruit immediately. In Fig. 10, we show one case that the robot first grasps the orange from the table and then places it in the target bin using the proposed framework. It should be noted that the purpose of this experiment is to illustrate that our framework can be potentially used for an integrated pick-and-place task. Therefore, our Fruit Classification model is not compared with other existing methods since it is not the focus of this work.

## V. CONCLUSIONS AND DISCUSSIONS

Our experiments demonstrate that the Transformer models can enable robotic grasping tasks in both the object classification and robot control domain. The results indicate that they outperform traditional models, such as CNN+LSTM, for classification tasks like slip detection and grasping outcome prediction. In addition, our Transformer-based grasping framework is able to select the grasping strength to safely grasp fruits with varying hardness and surface texture. We also visualize the attention flows of the Transformer models, which can potentially explain their effectiveness and efficacy. However, it is worth noting that the Transformer models are still a model-free method relying on the learned attention from rich data. Performance could be expected to be improved by incorporating model-based methods, such as physical contact models, as future work. Also, improving grasping task robustness and generalization via adversarially regularized policy learning [33] would be another direction to explore.

## REFERENCES

- [1] C. Blanes, M. Mellado, C. Ortiz, and Á. Valera, “Review. technologies for robot grippers in pick and place operations for fresh fruits and vegetables,” *Spanish Journal of Agricultural Research*, vol. 9, p. 1130.
- [2] Z. Xu, J. Wu, A. Zeng, J. B. Tenenbaum, and S. Song, “Densephysnet: Learning dense physical object representations via multi-step dynamic interactions,” in *Robotics: Science and Systems (RSS)*, 2019.
- [3] J. Wu, I. Yildirim, J. J. Lim, B. Freeman, and J. Tenenbaum, “Galileo: Perceiving physical object properties by integrating a physics engine with deep learning,” in *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [4] S. Luo, J. Bimbo, R. Dahiya, and H. Liu, “Robotic tactile perception of object properties: A review,” *Mechatronics*, vol. 48, pp. 54–67, 2017.
- [5] S. Wang, Y. She, B. Romero, and E. Adelson, “Gelsight wedge: Measuring high-resolution 3d contact geometry with a compact robot finger,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 6468–6475.
- [6] W. Yuan, S. Dong, and E. H. Adelson, “Gelsight: High-resolution robot tactile sensors for estimating geometry and force,” *Sensors*, vol. 17, no. 12, 2017.
- [7] N. Kuppawamy, A. Alspach, A. Uttamchandani, S. Creasey, T. Ikeda, and R. Tedrake, “Soft-bubble grippers for robust and perceptive manipulation,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 9917–9924.

- [8] 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 *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5633–5640.
- [9] J. Li, S. Dong, and E. Adelson, “Slip detection with combined tactile and visual information,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 7772–7777.
- [10] 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?” *arXiv preprint arXiv:1710.05512*, 2017.
- [11] G. Bertasius, H. Wang, and L. Torresani, “Is space-time attention all you need for video understanding?” in *ICML*, vol. 2, no. 3, 2021, p. 4.
- [12] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, and C. Schmid, “Vivit: A video vision transformer,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6836–6846.
- [13] J. Hughes, U. Culha, F. Giardina, F. Guenther, A. Rosendo, and F. Iida, “Soft manipulators and grippers: A review,” *Frontiers in Robotics and AI*, vol. 3, p. 69, 2016.
- [14] J. Jiang, G. Cao, A. Butterworth, T.-T. Do, and S. Luo, “Where shall i touch? vision-guided tactile poking for transparent object grasping,” *IEEE/ASME Transactions on Mechatronics*, 2022.
- [15] R. Calandra, A. Owens, D. Jayaraman, et al., “More than a feeling: Learning to grasp and regrasp using vision and touch,” *IEEE Robotics and Automation Letters*, vol. 3, pp. 3300–3307, 2018.
- [16] D. Kim, J. Lee, W.-Y. Chung, and J. Lee, “Artificial intelligence-based optimal grasping control,” *Sensors*, vol. 20, no. 21, 2020.
- [17] Y. Zhang, W. Yuan, Z. Kan, and M. Y. Wang, “Towards learning to detect and predict contact events on vision-based tactile sensors,” in *Conference on Robot Learning*, 2020, pp. 1395–1404.
- [18] S. Dong, D. Ma, E. Donlon, and A. Rodriguez, “Maintaining grasps within slipping bounds by monitoring incipient slip,” in *2019 International Conference on Robotics and Automation (ICRA)*, 2019.
- [19] N. Wettels, A. R. Parnandi, J.-H. Moon, G. E. Loeb, and G. S. Sukhatme, “Grip control using biomimetic tactile sensing systems,” *IEEE/ASME Transactions On Mechatronics*, vol. 14, no. 6, pp. 718–723, 2009.
- [20] J. Lin, R. Calandra, and S. Levine, “Learning to identify object instances by touch: Tactile recognition via multimodal matching,” in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 3644–3650.
- [21] W. Yuan, S. Wang, S. Dong, and E. Adelson, “Connecting look and feel: Associating the visual and tactile properties of physical materials,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5580–5588.
- [22] A. Vaswani, N. Shazeer, N. Parmar, et al., “Attention is all you need,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [23] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.
- [24] M. Shridhar, L. Manuelli, and D. Fox, “Perceiver-actor: A multi-task transformer for robotic manipulation,” *arXiv preprint arXiv:2209.05451*, 2022.
- [25] M. Monastirsky, O. Azulay, and A. Sintov, “Learning to throw with a handful of samples using decision transformers,” *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 576–583, 2023.
- [26] R. Yang, M. Zhang, N. Hansen, H. Xu, and X. Wang, “Learning vision-guided quadrupedal locomotion end-to-end with cross-modal transformers,” in *International Conference on Learning Representations*, 2022.
- [27] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [28] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, “Cable manipulation with a tactile-reactive gripper,” *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1385–1401, 2021.
- [29] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [30] S. Zuo, H. Jiang, Z. Li, T. Zhao, and H. Zha, “Transformer hawks process,” in *International conference on machine learning*. PMLR, 2020, pp. 11 692–11 702.
- [31] S. Abnar and W. Zuidema, “Quantifying attention flow in transformers,” *arXiv preprint arXiv:2005.00928*, 2020.
- [32] S. Hochreiter, Y. Bengio, et al., “Gradient flow in recurrent nets: the difficulty of learning long-term dependencies,” in *A Field Guide to Dynamical Recurrent Neural Networks*. IEEE Press, 2001.
- [33] Z. Zhao, S. Zuo, T. Zhao, and Y. Zhao, “Adversarially regularized policy learning guided by trajectory optimization,” *arXiv preprint arXiv:2109.07627*, 2021.