

# ViTacGen: Robotic Pushing with Vision-to-Touch Generation

Zhiyuan Wu<sup>1</sup>, Yijiong Lin<sup>2</sup>, Yongqiang Zhao<sup>1</sup>, Xuyang Zhang<sup>1</sup>, Zhuo Chen<sup>1</sup>, Nathan Lepora<sup>2</sup>, Shan Luo<sup>1</sup>

**Abstract**—Robotic pushing is a fundamental manipulation task that requires tactile feedback to capture subtle contact forces and dynamics between the end-effector and the object. However, real tactile sensors often face hardware limitations such as high costs and fragility, and deployment challenges involving calibration and variations between different sensors, while vision-only policies struggle with satisfactory performance. Inspired by humans’ ability to infer tactile states from vision, we propose ViTacGen, a novel robot manipulation framework designed for visual robotic pushing with vision-to-touch generation in reinforcement learning to eliminate the reliance on high-resolution real tactile sensors, enabling effective zero-shot deployment on visual-only robotic systems. Specifically, ViTacGen consists of an encoder-decoder vision-to-touch generation network that generates contact depth images, a standardized tactile representation, directly from visual image sequence, followed by a reinforcement learning policy that fuses visual-tactile data with contrastive learning based on visual and generated tactile observations. We validate the effectiveness of our approach in both simulation and real world experiments, demonstrating its superior performance and achieving a success rate of up to 86%. Code and data are available on <https://robot-perception-lab.github.io/vitacgen-website/>.

**Index Terms**—Force and Tactile Sensing, Reinforcement Learning, Deep Learning in Grasping and Manipulation

## I. INTRODUCTION

Robotic pushing is a fundamental manipulation task that involves applying forces to move objects toward a specified target region [1]. This task requires precise perception of the interactions between the robot and its environment during execution to enable accurate dynamic control [2]. In recent years, data-driven reinforcement learning (RL) approaches relying primarily on visual input have been widely explored for robotic pushing tasks. Some vision-based methods attempt to estimate contact dynamics. For instance, a nearby camera observing a compliant gripper can infer applied forces from the gripper’s deformation [3]. However, such approaches only provide indirect measurements of contact dynamics, which limits their accuracy and robustness. Therefore, researchers have introduced tactile sensing into robotic pushing tasks, leveraging its ability to provide detailed information about contact forces, surface properties, and interaction dynamics

[4], [5]. Compared to vision, tactile sensing provides localized and precise feedback [6], which is particularly important for capturing the fine-grained interactions required in robotic pushing tasks. By utilizing tactile modalities, these methods have demonstrated improved performance in capturing the complex relationships between the robot, object, and environment.

However, deploying tactile sensing in real world faces significant challenges. First, high-quality tactile sensors like vision-based sensors [7], [8] are inaccessible in many scenarios due to hardware limitations and specialized setups, and even when available, they typically require precise calibration [8] and remain prone to noise, wear, and physical damage during prolonged use [7]. Also, there exist manufacturing inconsistencies across different tactile sensors, which could cause subtle variations and manifest as noises, impairing model performance when deployed to new tactile sensors [9]. These limitations create a critical perception gap in tactile sensing deployment, and vision-only policies struggle to capture the subtle contact dynamics necessary for precise control, resulting in suboptimal performance in real-world pushing tasks.

Humans, in contrast, have the remarkable ability to infer tactile states from visual information [10]. Inspired by this, we propose ViTacGen, a novel robot manipulation framework designed for visual robotic pushing with vision-to-touch generation in reinforcement learning to eliminate the reliance on high-resolution real tactile sensors. Specifically, our method introduces two key components: an encoder-decoder vision-to-touch generation network **VT-Gen** that synthesizes tactile contact depth images [9], a standardized tactile representation, directly from visual image sequence, and an RL network **VT-Con** that learns robust policies through feature fusion and contrastive learning based on visual and generated tactile observations. Our ViTacGen is trained in simulation by two steps: (1) training VT-Gen to generate tactile contact depth images on paired visual and tactile data collected from expert trajectories by a pre-trained RL network [11] with visual and tactile observations, and (2) incorporating the frozen VT-Gen to train VT-Con based on visual and generated tactile observations. ViTacGen can perform zero-shot deployment on visual-only robotic systems, which eliminates the reliance on high-resolution real tactile sensors but still captures subtle dynamics and interactions between the end-effector and the object. Additionally, we solve the problem of manufacturing variations across different tactile sensors using contact depth maps as a standardized tactile representation. We conduct extensive experiments in both simulation and real world settings to validate the effectiveness of the proposed framework. The results demonstrate that ViTacGen achieves superior performance, achieving a success rate of up to 86% and outperforming baseline methods. Code and data will be open-

Manuscript received: May 19, 2025; Revised August 25, 2025; Accepted September 24, 2025.

This paper was recommended for publication by Editor Ashis Banerjee upon evaluation of the Associate Editor and Reviewers’ comments. This work was partially supported by the EPSRC project “ViTac: Visual-Tactile Synergy for Handling Flexible Materials” (EP/T033517/2). (Corresponding author: Shan Luo.)

<sup>1</sup>Zhiyuan Wu, Yongqiang Zhao, Xuyang Zhang, Zhuo Chen, and Shan Luo are with Department of Engineering, King’s College London, Strand, London, WC2R 2LS, United Kingdom, {zhiyuan.l.wu, yongqiang.zhao, xuyang.zhang, zhuo.7.chen, shan.luo}@kcl.ac.uk.

<sup>2</sup>Yijiong Lin and Nathan Lepora are with the Department of Engineering Mathematics and Bristol, Robotics Laboratory, University of Bristol, Bristol, BS8 1UB, United Kingdom, {yijiong.lin, n.lepora}@bristol.ac.uk.

Digital Object Identifier (DOI): see top of this page.

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

sourced once the paper is accepted.

Our main contributions are summarized as follows:

- We propose ViTacGen, a novel robot manipulation framework designed for visual robotic pushing with vision-to-touch generation in RL to eliminate the reliance on high-resolution real tactile sensors, enabling effective zero-shot deployment on visual-only robotic systems.
- We introduce an encoder-decoder vision-to-touch generation network VT-Gen that generates contact depth images as a standardized tactile representation directly from visual image sequence, with an RL network VT-Con that learns robust policies using feature fusion and contrastive learning on visual and generated tactile observations.
- We demonstrate the effectiveness of our proposed methods in both simulation and real world, through extensive qualitative and quantitative experiments, supported by comprehensive ablation studies.

## II. RELATED WORKS

### A. Visual-Tactile Representation Learning

Humans possess an extraordinary ability to perceive physical properties such as hardness, roughness, and texture through tactile sensing, which plays a critical role in providing feedback for tasks like grasping and manipulation. Research has increasingly focused on leveraging tactile information for robotic applications [6], with a significant body of work aiming to integrate tactile sensing into manipulation tasks, such as material classification and grasping [6], [12]. For example, tactile feedback has been widely utilized to improve grasping stability [13], which places high demands on the quality and robustness of tactile representations.

There is a natural synergy between visual and tactile modalities: vision provides a global understanding of objects, while touch offers detailed, localized information [11]. Combining these two modalities can yield richer and more comprehensive representations. With the rise of deep learning, recent research on visual-tactile fusion has shifted from simple data-level integration to joint representation learning. Luo *et al.* [14] pioneered the use of feature fusion for visual-tactile representations in cloth texture recognition. Subsequent studies [15], [16] have employed attention mechanisms to enhance the fusion of visual and tactile features, further improving performance in various tasks.

More recently, contrastive learning has emerged as a powerful paradigm, replacing traditional supervised approaches. Techniques such as Information Noise Contrastive Estimation (InfoNCE) [17] and Momentum Contrast (MoCo) [18] have been successfully applied to learn joint visual-tactile representations, as demonstrated in works like [19] and [20]. Additionally, in visual-tactile perception where both modalities are image-based, the structural similarity between visual and tactile images has been extensively leveraged to explore cross-modal generative modeling [21]. Generative models exploit this similarity to synthesize tactile data from visual input and infer tactile-related physical parameters such as friction [22] and material properties [19], particularly enabling vision-to-touch generation that allows the study of tactile properties without requiring direct physical interaction.

### B. Robotic Pushing

Robotic pushing refers to the task of manipulating objects by applying forces to push them into a specified target region [1]. It requires precise perception of the interactions between the robot and its environment during execution to enable accurate dynamic control [2]. While vision-based policies can estimate explicit contact dynamics, they rely heavily on precise visual perception such as end-effector deformation for contact parameter estimation [3]. To overcome these limitations, researchers have introduced tactile sensing into robotic pushing tasks, leveraging its ability to provide detailed information about contact forces, surface properties, and interaction dynamics [4], [5]. Lloyd *et al.* [4] first employed high-resolution optical tactile sensors in reactive robotic pushing to capture interaction details between end-effector and object. Lin *et al.* [9] and Yang *et al.* [5] pioneered in leveraging sim-to-real strategies with tactile sensing to train RL policies for robotic pushing. Recently, Lin *et al.* [23] investigated a set of bimanual pushing tasks tailored towards touch with RL. However, tactile sensing in real world applications depends on highly accurate vision-based tactile sensors [9], which can be inaccessible in many scenarios due to hardware limitations and specialized setups, and even when available, they typically require precise calibration [8] and remain prone to noise, wear, and physical damage during prolonged use [7]. Also, there exist manufacturing inconsistencies across different tactile sensors, which could cause subtle variations and manifest as noises, impairing model performance when deployed to new tactile sensors [9]. In this paper, we overcome these gaps by using generated contact depth maps from vision as a standardized tactile representation.

## III. PRELIMINARIES

### A. Reinforcement Learning

Reinforcement learning (RL) has been widely applied in robot manipulation tasks, enabling robots to learn complex behaviors through trial and error. Among various RL algorithms, the Soft Actor-Critic (SAC) algorithm [25] is a popular choice due to its ability to balance exploration and exploitation effectively. SAC is an off-policy actor-critic framework that incorporates entropy regularization into its objective function. This regularization encourages the policy to explore diverse actions, improving robustness and stability during training. Being off-policy, SAC leverages a replay buffer to learn from past experiences, which enhances sample efficiency [11]. The objective function  $J$  of SAC under policy  $\pi$  is defined as:

$$J(\pi) = \mathbb{E}_{(s_t, a_t) \sim \rho_\pi} \left[ \sum_t \gamma^t (R(s_t, a_t) + \alpha \mathcal{H}(\pi(\cdot | s_t))) \right], \quad (1)$$

where  $\mathbb{E}$  represents the expectation,  $s_t$  represents the environmental state at time step  $t$ , including the temporal window of visual and generated tactile contact depth observation sequences  $\mathcal{V}$  and  $\mathcal{C}$ , from time steps  $t-2$  to  $t$  (3 frames), as well as proprioception, *i.e.*, robot's tool center point (TCP) coordinates,  $a_t$  represents the robot's action,  $\rho_\pi$  represents the state-action distribution,  $\gamma$  represents the discount factor,  $R$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

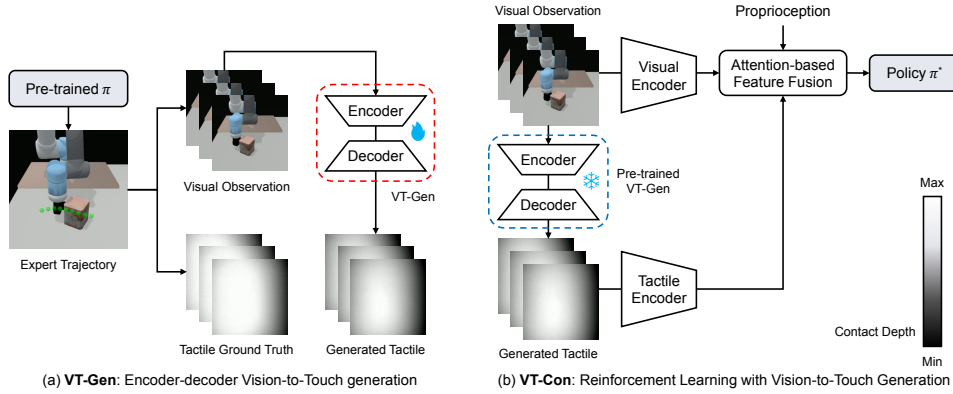


Fig. 1. The workflow of our proposed ViTacGen comprises two components: a VT-Gen for vision-to-touch generation, and a VT-Con for reinforcement learning on visual and generated tactile contact depth images with contrastive learning, and is trained in two stages: (a) In simulation, we collect paired visual and tactile data from expert trajectories by a pre-trained RL network [11] with visual and tactile observations, and use them to train a vision-to-touch generation network VT-Gen so as to generate contact depth images as a standardized tactile representation from visual input, and (b) we incorporate the frozen VT-Gen to train a goal-conditioned RL policy VT-Con based on proprioception *i.e.*, robot’s TCP coordinates, visual, and generated tactile observations. This architecture enables ViTacGen to perform zero-shot deployment in real world using only visual information during inference. Please note that we employ contact depth [24] to represent our tactile images following [9] (see Sec. III-B for details).

represents the reward function,  $\alpha$  represents the temperature parameter controlling the weight of the entropy term  $\mathcal{H}$ , and  $\mathcal{H}(\pi(\cdot|s_t))$  denotes the entropy of the policy  $\pi$  at state  $s_t$ :

$$\mathcal{H}(\pi(\cdot|s_t)) = -\mathbb{E}_{a_t \sim \pi(\cdot|s_t)} [\log \pi(a_t|s_t)]. \quad (2)$$

The reward function  $R(s_t, a_t) = -d_{\text{goal}} - d_{\text{TCP}}$  is designed as a dense reward that minimizes both the object-to-goal distance  $d_{\text{goal}}$  and the TCP-to-object distance  $d_{\text{TCP}}$ .

## B. Tactile Gym Simulator

We use Tactile Gym 2 [9] as our simulation platform, which provides efficient visual-tactile robot learning and manipulation capabilities through high-resolution optical tactile sensors. In Tactile Gym 2, the robotic pushing task is defined as moving the object along a desired 2D trajectory through a series of goal points. The observation space includes visual sensing, tactile sensing, and proprioception, *i.e.*, robot’s TCP coordinates. The principle of its tactile simulator [24] is to obtain contact depth images as a standardized tactile representation across different tactile sensors in our simulation environment. Based on PyBullet [26], it utilizes synthetic depth camera rendering to capture contact depth images within a virtual optical tactile sensor. When contacting with an object, the system computes the difference between the current depth image and a reference depth image taken when the sensor is not in contact to produce a penetration depth map that generalizes to arbitrary sensor shapes. To achieve computational efficiency for large-scale data generation, the simulator approximates the soft tip of the tactile sensor using rigid body physics, limiting contact stiffness and damping parameters during collision detection to allow objects to penetrate the simulated sensor tip in a manner that approximates the deformation of a real tip.

## IV. METHODOLOGY

### A. Overview

As illustrated in Fig. 1, our proposed ViTacGen framework comprises two key components: **VT-Gen** for encoder-decoder vision-to-touch generation and **VT-Con** to learn robust policies

through feature fusion and contrastive learning based on visual and generated tactile observations. The entire pipeline is correspondingly structured into two sequential phases to ensure effective training (simulation) and deployment (real world). Firstly, during training in simulation, we collect paired visual and tactile data from expert trajectories by a pre-trained RL network [11] with visual, tactile, and proprioception observations, and use them to train the vision-to-touch generation network VT-Gen to generate contact depth images from a visual image sequence. Secondly, we incorporate the frozen VT-Gen to train the RL policy VT-Con based on visual and generated tactile observations. This pipeline enables zero-shot transfer from simulated training to real world inference on visual-only robotic systems without real tactile sensors. The subsequent sections offer an in-depth analysis of both VT-Gen and VT-Con’s architectures.

### B. VTGen: Encoder-decoder Vision-to-touch Generation

Humans have the remarkable ability to infer tactile states from visual information [10]. Inspired by this, we propose a VT-Gen for encoder-decoder vision-to-touch generation to infer the tactile sensation of the contact region directly from a image sequence of visual input. As depicted in Fig. 1 (a), given a visual image sequence  $\mathcal{V} = \{v_1, \dots, v_N\}$ , it is processed by VT-Gen to generate a current tactile contact depth image  $c^{\text{gen}}$ , which is then repeated  $N$  times to form a generated tactile contact depth sequence  $\mathcal{C}$  with  $N$  frames to align with the length of the visual sequence, serving as the tactile observation for VT-Con.

Specifically, we first concatenate  $\mathcal{V}$  along the channel dimension and extract features with coarse-to-refine encoding, so as to capture hierarchical and multi-scale visual features at multi scales [27]. We pass  $\mathcal{V}$  through a coarse encoder  $\mathcal{E}_{\text{coarse}}$ , which extracts the initial visual feature map  $f^{\text{coarse},v}$ . This is followed by a cross-modal attention operation  $\mathcal{A}_{\text{cm}}$  [28], formulated as:

$$f_{\text{cm}}^v = \mathcal{C}[\mathcal{A}_1^{\text{cm}}(f^{\text{coarse},v}, p), \dots, \mathcal{A}_h^{\text{cm}}(f^{\text{coarse},v}, p)]w_0, \quad (3)$$

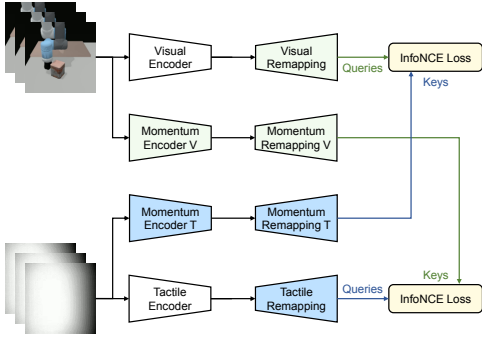


Fig. 2. To better fuse the encoded visual and tactile features, we optimize the visual and tactile encoders in VT-Con (Fig. 1 (b)) with MoCo [18] for contrastive learning between visual and tactile modalities. We utilize two momentum-based encoders to obtain momentum visual and tactile features, which are then aligned with the original features of the other modality through the InfoNCE loss [17].

where  $p$  represents a learnable positional embedding [29],  $f_{cm}^v$  refers to the processed visual feature map,  $h$  is the number of attention heads that empirically set to 8, and  $w_0$  represents the weight matrix used for output. The cross-modal attention operation  $\mathcal{A}_i^{cm}$  is defined as:

$$\mathcal{A}_{cm}(x, y) = \text{softmax}\left(\frac{qk^T}{\sqrt{d}}\right) \cdot v, \quad (4)$$

with

$$q = w_q \cdot x, \quad k = w_k \cdot y, \quad v = w_v \cdot y, \quad (5)$$

where  $w$  are learnable projection matrices [28].  $f_{cm}^v$  is then processed by a refine encoder  $\mathcal{E}_{refine}$ , to produce a deeper and more refined visual feature map  $f_{refine}^v$ . Next,  $f_{refine}^v$  is processed through a series of identical residual blocks [30], which enhance feature representation while preserving spatial information. Finally, a hierarchical decoder, composed of transposed convolutional layers and upsampling operations, generates the final contact depth prediction  $c^{gen}$ .

### C. VT-Con: RL with Visual-Tactile Contrastive Learning

1) *Feature Extraction*: As shown in Fig. 1 (b), we consider the visual sequence  $\mathcal{V} = \{v_1, \dots, v_N\}$  for visual observation, and the generated tactile sequence  $\mathcal{C} = \{c_1, \dots, c_N\}$  for tactile observation, where  $N$  denotes the number of frames in the visual and tactile sequences, respectively, empirically set to 3 for capturing the velocity and acceleration information. Each element  $v_i$  or  $c_i$  is represented as a tensor resized to the same dimension  $\mathbb{R}^{H \times W \times C}$ , where  $H$ ,  $W$ , and  $C$  correspond to the height, width, and channel count of the frames. To extract features from the visual and tactile inputs, we first concatenate  $\mathcal{V}$  and  $\mathcal{C}$  respectively along their channel dimensions. Then, we employ two structurally identical Convolutional Neural Networks (CNNs)  $\mathcal{E}^v$  and  $\mathcal{E}^c$  as encoders for the visual and tactile modalities, resulting in a visual feature map  $f^v$  and a tactile feature map  $f^c$ .

2) *Contrastive Learning between Vision and Touch*: In order to more efficiently fuse the encoded visual and tactile features  $f^v$  and  $f^c$ , we optimize  $\mathcal{E}^v$  and  $\mathcal{E}^c$  with Momentum Contrast (MoCo) [18] for contrastive learning between vision and touch, inspired by [11]. As shown in Fig. 2, we

utilize two momentum-based encoders  $\mathcal{M}^v$  and  $\mathcal{M}^c$  to obtain momentum visual and tactile features  $m^v$  and  $m^c$  from  $\mathcal{V}$  and  $\mathcal{C}$ , respectively. These momentum encoders are structurally identical to the original encoders  $\mathcal{E}^v$  and  $\mathcal{E}^c$  but are updated using a momentum-based approach, formulated as:

$$\mathcal{M}^v \leftarrow \eta \mathcal{M}^v + (1 - \eta) \mathcal{E}^v, \quad \mathcal{M}^c \leftarrow \eta \mathcal{M}^c + (1 - \eta) \mathcal{E}^c, \quad (6)$$

where  $\eta \in [0, 1)$  is the momentum coefficient that controls the update rate. We compute the contrastive loss between the visual and tactile modalities using the InfoNCE loss [17], aligning  $f^v$  with  $m^c$  and  $f^c$  with  $m^v$ . (see Sec. IV-D for details). This contrastive learning module is integrated into VT-Con's RL process by incorporating the contrastive loss into the actor's training. Specifically, after each time the RL updates, the contrastive loss is computed, and a back-propagation is performed to optimize the encoders  $\mathcal{E}^v$  and  $\mathcal{E}^c$ .

3) *Attention-based Visual-Tactile Feature Fusion*: We introduce an attention-based visual-tactile feature fusion operation [28], similar to Eq. 3, to integrate  $f^v$  and  $f^c$ , which can be formulated as follows:

$$f^{fuse} = C[\mathcal{A}_1^{cm}(f^v, f^c), \dots, \mathcal{A}_h^{cm}(f^v, f^c)]w_0, \quad (7)$$

where  $f^{fuse}$  refers to the fused feature map. Along with the fused feature map  $f^{fuse}$ , the flattened and MLP-processed proprioception, *i.e.*, robot's TCP coordinates, are concatenated to form the complete observation vector for robot reinforcement learning and can be compatible with various RL algorithms, such as SAC and PPO.

### D. Loss Function

1) *Loss Function for VT-Gen*: In VT-Gen, we utilize the VGG loss [31] to supervise the training of our vision-to-touch generation. The VGG loss is designed to evaluate the perceptual similarity between the generated contact depth image  $c^{gen}$  and the ground truth contact depth image  $c^{gt}$  by comparing their high-level feature representations extracted from a pre-trained VGG network. Specifically, the loss  $\mathcal{L}^{vgg}$  is defined as:

$$\mathcal{L}^{vgg} = \|\phi(c^{gen}) - \phi(c^{gt})\|^2, \quad (8)$$

where  $\phi$  denotes the feature extraction process using the encoder of a pre-trained VGG network [31],  $c^{gt}$  represents the ground truth contact depth image, and  $\|\cdot\|^2$  corresponds to the mean squared error.

2) *Loss Function for VT-Con*: In VT-Con, we align visual and tactile modalities using InfoNCE loss [17] as the contrastive loss  $\mathcal{L}^c$ , which consists of a vision-to-touch loss  $\mathcal{L}^{vt}$  and a touch-to-vision loss  $\mathcal{L}^{tv}$ , defined as:

$$\mathcal{L}^{vt} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp\left(\frac{f_i^v \cdot m_i^c}{\tau}\right)}{\sum_{\substack{j=1 \\ j \neq i}}^B \exp\left(\frac{f_i^v \cdot m_j^c}{\tau}\right)}, \quad (9)$$

and

$$\mathcal{L}^{tv} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp\left(\frac{f_i^c \cdot m_i^v}{\tau}\right)}{\sum_{\substack{j=1 \\ j \neq i}}^B \exp\left(\frac{f_i^c \cdot m_j^v}{\tau}\right)}, \quad (10)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.



Fig. 3. Real world objects used in our experiments: (a) tea box, (b) meat can, (c) mug, and (d) unseen objects for zero-shot test including i) olive jar, ii) apple, iii) coffee can, iv) soup can, and v) ceramic cup.

where  $B$  denotes the number of samples (or batch size in implementation), and  $\tau$  is the temperature parameter, empirically set to 0.1. The overall contrastive loss  $\mathcal{L}^c$  is then expressed as:

$$\mathcal{L}^{con} = \mathcal{L}^{vt} + \mathcal{L}^{tv}. \quad (11)$$

This loss enables our VT-Con to optimize  $\mathcal{E}^v$  and  $\mathcal{E}^c$  and enhance the multi-modal understanding capability during the RL training, fusing visual and tactile features more efficiently.

## V. EXPERIMENTS

### A. Experimental Setup

To validate the effectiveness of ViTacGen, we conducted comprehensive evaluations in both simulated and real world environments. Our simulation experiments were performed using Tactile Gym 2 [9] (see Sec. III-B for details). We selected three representative objects: a tea box, a potted meat box, and a mug from the YCB dataset [32], whose digitally sampled meshes were incorporated into our simulation environment. We also selected five real-world unseen objects for testing. The objects are shown in Fig. 3. We employ a UR5e robot arm equipped with a pushing end-effector to manipulate objects with a control frequency of 500Hz, with each manipulation episode limited to 350 timesteps within a constrained workspace of  $800 \times 600$  mm in xy-plane at a fixed height of 2 cm above the table surface. We employ an external Intel RealSense D435 RGB camera with a field of view (FOV) of  $42^\circ$ , mounted 1 meter away from the robot base center with a  $30^\circ$  downward viewing angle, and the camera parameters are kept consistent in the corresponding simulation training. Our simulation experiments adhered to the standardized configurations of Tactile Gym 2 with a maximum of 350 timesteps per episode, and we implemented object mass, surface friction coefficient, and visual domain randomization including camera view, lighting, background, and color variations [33] to improve our model’s adaptability and robustness. The pushing trajectories are randomly generated using OpenSimplex noise.

For our VT-Gen, we collect visual and tactile data from 1,000 manipulation sequences and train our model for 200 epochs with a batch size of 64, splitting the data in a ratio of 7:2:1 for training, validation, and testing. For VT-Con, we employ the Soft Actor-Critic (SAC) [25] algorithm from Stable Baselines 3 [34] as our RL backbone, with a max episode length of 350. We train our model for 1,000,000 timesteps with a batch size of 64 and a buffer size of 20,000 for experience replay. Both visual and tactile images are resized to

TABLE I  
QUANTITATIVE RESULTS FOR VISION-TO-TOUCH GENERATION IN RL SETTING. WE REPORT PSNR, SSIM [36], AND LPIPS [37].

Object Type	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Tea Box	30.75	0.9482	0.0101
Meat Box	20.50	0.8657	0.0327
Mug	20.25	0.8222	0.0417

$128 \times 128$  pixels. The success termination criterion is defined as the distance between the object center and the goal center being less than 2.5 cm. For optimizer we employ Adam [35], with a learning rate of  $1e-4$ , and an epsilon value of  $1e-8$ .

### B. Evaluation Metrics and Baselines

To evaluate the performance of our ViTacGen in both simulation and real world scenarios, we employ several key metrics. In the simulation environment, we assess cumulative rewards and episode length to evaluate the effectiveness of our learning algorithm. For both simulation and real-world experiments, we measure the final distance error, which quantifies the accuracy of the robot’s manipulation tasks. Additionally, we calculate the success rates based on this distance error, with a threshold of less than 2.5 cm. For baseline comparison, in simulation we compare our method with visual-only policy, tactile-only policy, and visual & tactile policy in [9]. We evaluate the performance of our ViTacGen with 1) visual-only input and 2) visual & actual tactile input. The latter uses attention and contrastive learning, but without vision-to-touch generation. In real world, we compare our method with visual-only policy as we deploy our ViTacGen in a visual-only robotic system.

### C. Effectiveness of Vision-to-Touch Generation

To begin with, we quantitatively evaluate the effectiveness of our VT-Gen using standard metrics including PSNR, SSIM [36], and LPIPS [37]. As shown in Tab. I, VT-Gen achieves great performance on generating tactile contact depth map from visual image sequence, as evidenced by the tactile contact depth maps shown in Fig. 4 (ii) and (iii). Its robust performance can be attributed to its coarse-to-refine cross-modal based encoders, which capture hierarchical and multi-scale visual features. The high generation quality achieved by ViTacGen with visual-only input also demonstrates its effectiveness in maintaining manipulation performance. Notably, our VT-Gen features a modest model size of 146.74 MiB and achieves an inference speed of 305.90 FPS, ensuring real-time operation and practical deployability.

### D. Simulation Robotic Pushing Results

1) *Quantitative Results*: As demonstrated in the quantitative results given in Tab. II, simply employing single modality (visual or tactile) fails to learn a good strategy for robotic pushing, as vision only provides a rough estimation of the scene, while tactile sensing only focuses on the interactions between end-effectors and objects, leading to limited perception capability. The Visual & Tactile baseline shows only marginal improvement because it simply concatenates visual and tactile features without effective cross-modal alignment,

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE II

QUANTITATIVE RESULTS FOR ROBOTIC PUSHING IN SIMULATION. WE COMPARE TWO VERSIONS OF OUR ViTACGEN FRAMEWORK: ONE WITH VISUAL-ONLY INPUTS AND THE OTHER WITH VISUAL & ACTUAL TACTILE DATA FOR REFERENCE. THESE ARE EVALUATED AGAINST BASELINES INCLUDING VISUAL-ONLY, TACTILE-ONLY, AND COMBINED VISUAL-TACTILE APPROACHES [9]. WE CONDUCT 100 TRIALS FOR EACH EXPERIMENTAL CONDITION. WE REPORT CUMULATIVE REWARDS, EPISODE LENGTH, DISTANCE ERROR, AND SUCCESS RATE WITH A THRESHOLD OF 2.5 CM. THE BEST RESULTS ARE HIGHLIGHTED IN RED, WHILE THE SECOND-BEST RESULTS ARE HIGHLIGHTED IN BLUE.

Object Type	Method	Rewards (mean±std) ↑	Epi. Len. (mean±std) ↓	Dist. Err. (mm) ↓	Succ. Rate (%) ↑
Tea Box	Visual only	-155.73±63.06	335.93±38.15	47.54	12.0
	Tactile only	-150.10±62.03	337.10±36.72	50.62	11.0
	Visual & Tactile	-147.01±72.21	334.80±39.34	50.02	13.0
	<b>Ours</b> (Visual only)	<b>-84.35±79.03</b>	<b>267.24±51.90</b>	<b>27.81</b>	<b>84.0</b>
	<b>Ours</b> (Visual & Tactile)	<b>-44.83±22.64</b>	<b>268.57±30.20</b>	<b>23.08</b>	<b>92.0</b>
Meat Can	Visual only	-169.27±80.65	330.02±45.84	62.23	16.0
	Tactile only	-149.75±17.06	324.44±45.39	53.77	25.0
	Visual & Tactile	-125.72±79.06	313.98±55.18	53.75	30.0
	<b>Ours</b> (Visual only)	<b>-88.19±79.86</b>	<b>250.90±48.50</b>	<b>29.81</b>	<b>81.0</b>
	<b>Ours</b> (Visual & Tactile)	<b>-44.49±20.96</b>	<b>251.66±41.86</b>	<b>30.38</b>	<b>86.0</b>
Mug	Visual only	-112.65±52.30	328.78±43.09	54.85	20.0
	Tactile only	-106.35±50.17	317.51±49.73	47.19	31.0
	Visual & Tactile	-106.01±50.06	311.20±51.95	42.83	38.0
	<b>Ours</b> (Visual only)	<b>-41.53±19.72</b>	<b>266.03±43.91</b>	<b>27.39</b>	<b>86.0</b>
	<b>Ours</b> (Visual & Tactile)	<b>-34.92±13.86</b>	<b>270.99±38.21</b>	<b>27.07</b>	<b>83.0</b>

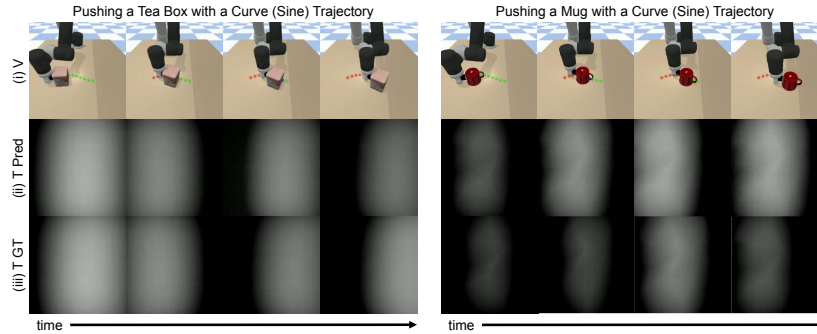


Fig. 4. Qualitative results of our ViTACGen in simulation, where we present a tea box and a mug with a curve (sine) trajectories. (i) Visual input with visualized pushing trajectories, where the green points refer to the target trajectory and the red points refer to the successfully executed object trajectory (ii) tactile contact depth prediction generated by VT-Gen (iii) tactile observation. Tactile contact depth maps correspond to frontal end-effector observation.

which hinders convergence during RL training. Without contrastive learning, the model is unable to establish meaningful relationships between the modalities, resulting in ineffective learning. In contrast, our VT-Con employs contrastive learning and attention-based fusion to properly couple cross-modal representations, facilitating convergence and leading to significant performance improvements across all metrics. Compared to the baseline, ViTACGen with visual and actual tactile data shows an improvement of 64.61% to 69.50% on cumulative rewards, 12.92% to 23.30% on episode length, and 13.36 mm to 26.94 mm on distance error, and achieves a success rate of 83% to 94%. Furthermore, our ViTACGen introduces VT-Gen to eliminate the dependency on tactile modality, instead predicting tactile contact depth images through visual observation. As shown in the table, even with visual modality only as input, ViTACGen maintains good performance across all metrics, significantly outperforming even the baseline with visual and tactile observations, and in some metrics even surpassing its own performance with visual and tactile observations, such as episode length for tea box, meat can, and mug, and the distance error for meat can. It maintains a high success rate of 80% to 86%, demonstrating our effectiveness and robustness.

2) *Qualitative Results:* We provide qualitative simulation results of our ViTACGen in Fig. 4, where pushing trajectories are visualized in visual images for reference. It can be found in Fig. 4 (i) that ViTACGen demonstrates robust and stable

TABLE III

QUANTITATIVE RESULTS FOR ROBOTIC PUSHING IN REAL WORLD. SINCE OUR ViTACGEN IN REAL WORLD SETTING IS DEPLOYED ON A VISUAL-ONLY ROBOTIC SYSTEM, WE COMPARE IT AGAINST VISUAL-ONLY INPUT BASELINE IN [9]. WE CONDUCT 50 TRIALS FOR EACH EXPERIMENTAL CONDITION AND MANUALLY MEASURE AVERAGE DISTANCE ERROR AND SUCCESS RATE WITH A THRESHOLD OF 2.5 CM.

Object Type	Method	Dist. Err. (cm) ↓	Succ. Rate (%) ↑
Tea Box	baseline	6.5±1.9	14.0
	<b>Ours</b>	<b>2.6±0.8</b>	<b>76.0</b>
Meat Can	baseline	8.2±1.9	8.0
	<b>Ours</b>	<b>1.9±0.5</b>	<b>82.0</b>
Mug	baseline	7.2±2.1	10.0
	<b>Ours</b>	<b>1.8±0.7</b>	<b>86.0</b>

performance across objects of various shapes. We also present the tactile prediction by VT-Gen and the corresponding tactile observation in Fig. 4 (ii) and (iii).

### E. Real World Robotic Pushing Demonstration

The performance of ViTACGen in real world scenarios is evaluated both quantitatively and qualitatively, as presented in Tab. III and Fig. 5, respectively. We conduct 50 trials for each experimental condition and manually measure average distance error and success rate. In Fig. 5, we visualize the trajectories in simulation and their projections onto both simulated and real world environments to demonstrate the system's effectiveness. The experimental results indicate that ViTACGen achieves robust zero-shot transfer to real world

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

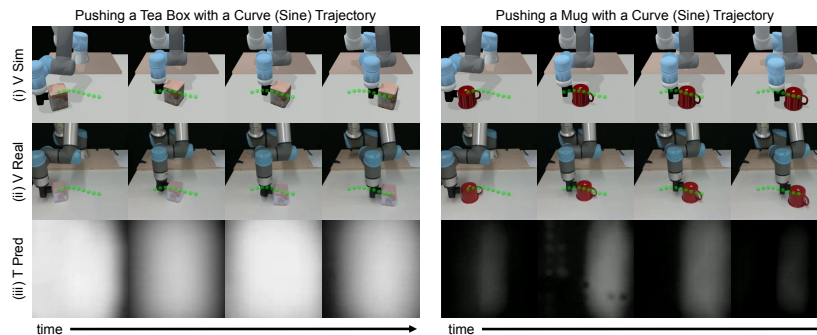


Fig. 5. Demonstration of ViTacGen in real world scenarios, where we present a tea box and a mug with a curve (sine) trajectories. The trajectories are visualized in simulation and projected onto both simulated and real world environments to demonstrate the system’s performance. (i) Simulation visual sequences for reference (ii) real world visual sequences (iii) tactile contact depth images generated from real world visual input, corresponding to frontal end-effector observation. More qualitative results are available in our supplementary material.

TABLE IV

QUANTITATIVE RESULTS FOR PUSHING UNSEEN OBJECTS IN REAL WORLD. WE CONDUCT 50 TRIALS FOR EACH EXPERIMENTAL CONDITION AND MANUALLY MEASURE AVERAGE DISTANCE ERROR AND SUCCESS RATE WITH A THRESHOLD OF 4.0 CM.

Object Type	Dist. Err. (cm) ↓	Succ. Rate (%) ↑
Apple	$3.3 \pm 1.1$	82.0
Coffee Box	$4.1 \pm 2.0$	72.0
Ceramic Cup	$3.9 \pm 1.8$	74.0
Olive Jar	$3.7 \pm 1.7$	78.0
Soup Can	$4.1 \pm 1.7$	70.0

scenarios, primarily due to VT-Gen’s capability in capturing subtle contact forces and dynamics between the end-effector and the object. Although the tactile predictions for real world scenarios exhibit some noise artifacts (Fig.5 (iii)), attributable to the sim-to-real discrepancy in visual modalities, as evident in Fig.5 (i) and (ii), these perturbations do not significantly impact the RL algorithm’s execution. Notably, the distance errors in real world experiments are consistently lower than their simulation counterparts, which can be attributed to our implementation of a threshold-based termination criterion in the simulation environment. We also provide real world robotic pushing results on unseen objects in Tab. IV, where our ViTacGen demonstrates good generalizability and robustness in handling unseen objects with diverse shapes. On average, our method achieves 75.2% success rate with 3.8 cm distance error across these diverse unseen objects, with the best performance on the apple, which is the smallest object among all the objects, which is easiest to control during pushing, and the worst on the soup can, likely due to its high center of gravity and tendency to spin during pushing motions. More qualitative results are available in our demo video.

#### F. Ablation Study

We conduct ablation studies on fusion module and contrastive learning module. The results are shown in Tabs. V and VI. For fusion module, we compare addition, concatenation, and attention [28], demonstrating that attention-based fusion achieves the best performance. For the contrastive learning module, we compare SimCLR [38] and MoCo [18], showing that MoCo outperforms SimCLR across all metrics.

## VI. CONCLUSION

This paper presents ViTacGen, a novel robot manipulation framework designed for visual robotic pushing with vision-to-

touch generation in reinforcement learning to eliminate the reliance on high-resolution real tactile sensors, inspired by human’s remarkable ability of predicting tactile states from vision to optimize manipulation. Specifically, we introduce an encoder-decoder vision-to-touch generation network that generates contact depth images, a standardized tactile representation, directly from visual image sequence, followed by a reinforcement learning policy that fuses visual-tactile data with contrastive learning based on visual and generated tactile observations. ViTacGen enables zero-shot deployment on visual-only robotic systems. We conduct extensive evaluation in both simulated and real world environments, demonstrating the effectiveness of our proposed methods.

While our ViTacGen demonstrates effectiveness, it contains certain limitations. Our method employs contact depth as a high-level tactile representation to provide a perception for contact geometry rather than generating precise tactile images, which reduces the sim-to-real gap but limits our ability to perceive detailed physical properties such as local force distributions and fine-grained surface geometry. Future work could focus on explicitly predicting these physical properties to enhance policy performance. Additional promising research direction involves exploring methods to infer shear forces [39] applied on tactile sensors directly from visual observations of object motions, which would further bridge the gap between visual and tactile modalities.

## ACKNOWLEDGMENT

The authors would like to thank Yupeng Wang for his help in hardware support for the experiments.

## REFERENCES

- [1] B. Deng, Y. Lin, M. Yang, and N. F. Lepora, “Coarse-to-fine robotic pushing using touch, vision and proprioception,” *IEEE Robotics and Automation Letters*, 2024.
- [2] K.-T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez, “More than a million ways to be pushed. a high-fidelity experimental dataset of planar pushing,” in *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2016, pp. 30–37.
- [3] Y. Zhu, M. Hao, X. Zhu, Q. Bateux, A. Wong, and A. M. Dollar, “Forces for free: Vision-based contact force estimation with a compliant hand,” *Science Robotics*, vol. 10, no. 103, p. eadq5046, 2025.
- [4] J. Lloyd and N. F. Lepora, “Goal-driven robotic pushing using tactile and proprioceptive feedback,” *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 1201–1212, 2021.

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE V

ABLATION STUDY ON THE FUSION MODULE. WE CONDUCT 100 TRIALS FOR EACH EXPERIMENTAL CONDITION. WE REPORT CUMULATIVE REWARDS, EPISODE LENGTH, DISTANCE ERROR, AND SUCCESS RATE WITH A THRESHOLD OF 4.0 CM.

Fusion Module	Rewards (mean±std) ↑	Epi. Len. (mean±std) ↓	Dist. Err. (mm) ↓	Succ. Rate (%) ↑
Addition	-70.00±36.47	265.01±47.70	39.93	77.0
Concatenation	-67.84±59.26	251.81±35.21	36.37	90.0
Attention [28]	<b>-66.92±55.28</b>	<b>247.82±27.81</b>	<b>35.89</b>	<b>95.0</b>

TABLE VI

ABLATION STUDY ON THE CONTRASTIVE LEARNING MODULE. WE CONDUCT 100 TRIALS FOR EACH EXPERIMENTAL CONDITION. WE REPORT CUMULATIVE REWARDS, EPISODE LENGTH, DISTANCE ERROR, AND SUCCESS RATE WITH A THRESHOLD OF 4.0 CM.

Fusion Module	Rewards (mean±std) ↑	Epi. Len. (mean±std) ↓	Dist. Err. (mm) ↓	Succ. Rate (%) ↑
SimCLR [38]	-70.03±49.30	251.97±36.26	37.58	89.0
MoCo [18]	<b>-66.92±55.28</b>	<b>247.82±27.81</b>	<b>35.89</b>	<b>95.0</b>

- [5] M. Yang, Y. Lin, A. Church, J. Lloyd, D. Zhang, D. A. Barton, and N. F. Lepora, "Sim-to-real model-based and model-free deep reinforcement learning for tactile pushing," *IEEE Robotics and Automation Letters*, 2023.
- [6] S. Luo, J. Binbo, R. Dahiya, and H. Liu, "Robotic tactile perception of object properties: A review," *Mechatronics*, vol. 48, pp. 54–67, 2017.
- [7] 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.
- [8] 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 *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1927–1934.
- [9] Y. Lin, J. Lloyd, A. Church, and N. F. Lepora, "Tactile gym 2.0: Sim-to-real deep reinforcement learning for comparing low-cost high-resolution robot touch," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10754–10761, 2022.
- [10] F. N. Newell, A. T. Woods, M. Mernagh, and H. H. Bühlhoff, "Visual, haptic and crossmodal recognition of scenes," *Experimental Brain Research*, vol. 161, no. 2, pp. 233–242, 2005.
- [11] F. Lygerakis, V. Dave, and E. Rueckert, "M2curl: Sample-efficient multimodal reinforcement learning via self-supervised representation learning for robotic manipulation," *arXiv preprint arXiv:2401.17032*, 2024.
- [12] G. Cao, Y. Zhou, D. Bollegala, and S. Luo, "Spatio-temporal attention model for tactile texture recognition," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9896–9902.
- [13] Z. Deng, Y. Jonetzko, L. Zhang, and J. Zhang, "Grasping force control of multi-fingered robotic hands through tactile sensing for object stabilization," *Sensors*, vol. 20, no. 4, p. 1050, 2020.
- [14] S. Luo, W. Yuan, E. Adelson, A. G. Cohn, and R. Fuentes, "Vitic: Feature sharing between vision and tactile sensing for cloth texture recognition," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2722–2727.
- [15] S. Cui, R. Wang, J. Wei, J. Hu, and S. Wang, "Self-attention based visual-tactile fusion learning for predicting grasp outcomes," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5827–5834, 2020.
- [16] Y. Chen, M. Van der Merwe, A. Sipos, and N. Fazeli, "Visuo-tactile transformers for manipulation," in *6th Annual Conference on Robot Learning*, 2022.
- [17] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [18] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 9729–9738.
- [19] F. Yang, C. Ma, J. Zhang, J. Zhu, W. Yuan, and A. Owens, "Touch and go: Learning from human-collected vision and touch," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 35, pp. 8081–8103, 2022.
- [20] Z. Wu, Y. Zhao, and S. Luo, "Convitic: Aligning visual-tactile fusion with contrastive representations," *arXiv preprint arXiv:2506.20757*, 2025.
- [21] J.-T. Lee, D. Bollegala, and S. Luo, "touching to see" and "seeing to feel": Robotic cross-modal sensory data generation for visual-tactile perception," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 4276–4282.
- [22] G. Cao, J. Jiang, N. Mao, D. Bollegala, M. Li, and S. Luo, "Vis2hap: Vision-based haptic rendering by cross-modal generation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 12443–12449.
- [23] Y. Lin, A. Church, M. Yang, H. Li, J. Lloyd, D. Zhang, and N. F. Lepora, "Bi-touch: Bimanual tactile manipulation with sim-to-real deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 8, no. 9, pp. 5472–5479, 2023.
- [24] A. Church, J. Lloyd, N. F. Lepora *et al.*, "Tactile sim-to-real policy transfer via real-to-sim image translation," in *Conference on Robot Learning*. PMLR, 2022, pp. 1645–1654.
- [25] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.
- [26] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.
- [27] Z. Wu, Y. Feng, C.-W. Liu, F. Yu, Q. Chen, and R. Fan, "S<sup>3</sup>m-net: Joint learning of semantic segmentation and stereo matching for autonomous driving," *IEEE Transactions on Intelligent Vehicles*, 2024.
- [28] 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," in *International Conference on Learning Representations (ICLR)*, 2020.
- [29] Z. Wu, X. Song, S. Wang, W. Liu, J. Yang, Z. Cheng, S. Chen, T. Shang, W. Sun, S. Luo *et al.*, "Cdi3d: Cross-guided dense-view interpolation for 3d reconstruction," *arXiv preprint arXiv:2503.08005*, 2025.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [31] K. Simonyan, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [32] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *2015 international conference on advanced robotics (ICAR)*. IEEE, 2015, pp. 510–517.
- [33] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.
- [34] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," *Journal of Machine Learning Research*, vol. 22, no. 268, pp. 1–8, 2021.
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [36] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [37] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 586–595.
- [38] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [39] Y. Lin, J. Huang, M. Zimmer, Y. Guan, J. Rojas, and P. Weng, "Invariant transform experience replay: Data augmentation for deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6615–6622, 2020.