

ViTaS: Visual Tactile Soft Fusion Contrastive Learning for Visuomotor Learning




Fig. 1: ViTaS is capable of handling various simulation and real-world manipulation tasks including transparent objects and self-occluded scenarios, by fusing visual and tactile features for effective policy learning.

Abstract—Tactile information plays a crucial role in human manipulation tasks and has recently garnered increasing attention in robotic manipulation. However, existing approaches mostly focus on the alignment of visual and tactile features and the integration mechanism tends to be direct concatenation. Consequently, they struggle to effectively cope with occluded scenarios due to neglecting the inherent complementary nature of both modalities and the alignment may not be exploited enough, limiting the potential of their real-world deployment. In this paper, we present ViTaS, a simple yet effective framework that incorporates both visual and tactile information to guide the behavior of an agent. We introduce *Soft Fusion Contrastive Learning*, an advanced version of conventional contrastive learning method and a CVAE module to utilize the alignment and complementarity within visuo-tactile representations. We demonstrate the effectiveness of our method in 12 simulated and 3 real-world environments, and our experiments show that ViTaS significantly outperforms existing baselines. We are committed to making the code open-source.

I. INTRODUCTION

Humans are adept at performing complex manipulation tasks, such as spinning an object or cleaning a table. While vision plays a critical role, other modalities, particularly touch, also provide rich information for these activities. Interestingly, visual and tactile information often exhibit significant relevance and complementarity [1]. For individuals with visual impairments, a clearer mental reconstruction of an original visual image can be achieved by combining a blurred visual perception with tactile information [2].

Most previous visuomotor learning algorithms have relied primarily on visual information to address manipulation

tasks [3]–[7]. Recently, several efforts have aimed to incorporate tactile information to improve the performance of algorithms. M3L [8] directly concatenates visual and tactile inputs and feeds them into MAE, while VTT [9] segments visual and tactile data into patches and uses a transformer to extract representations. However, these approaches fail to exploit the inherent correspondence between visual and tactile information, resulting in suboptimal performance on challenging manipulation tasks. Moreover, many prior approaches overlook the complementary nature of visual and tactile modalities, which could undermine performance in self-occluded scenarios and potentially limit their practical deployment. Given these limitations, we pose the question: *how can we more effectively fuse visual and tactile information to enhance the performance of visuomotor algorithms?*

Drawing on prior research in human physiology regarding the processing of visuo-tactile information, we propose **Visual Tactile Soft Fusion Contrastive Learning** (ViTaS), a novel visuo-tactile representation learning framework for visuomotor learning. Generally, ViTaS can be divided into two parts. Firstly, given the inherent relevance between visual and tactile modalities, we utilize contrastive learning to align the embeddings of visual data with their corresponding tactile information in the latent space. Notably, we employ *soft fusion contrastive learning*, a novel method extending the RGB single-modality framework presented in [10] to fuse features in alternating modalities. Second, inspired by humans’ remarkable ability to reconstruct clear mental images from blurred visual input when supplemented with




Fig. 2: Method overview. ViTaS takes vision and touch as inputs, which are then processed through separate CNN encoders. Encoded embeddings are utilized by *soft fusion contrastive* approach, yielding fused feature representation for policy network. A CVAE-based reconstruction framework is also applied for cross-modal integration.

tactile information, we integrate a conditional variational autoencoder (CVAE) [11] that leverages the complementary nature of two modalities. The CVAE reconstructs visual observations from fused visuo-tactile embeddings, enforcing cross-modal consistency and enhancing the quality of multimodal representations.

To evaluate the performance of our algorithm, we conduct extensive experiments across both simulated and real-world domains with RL and IL paradigms. Our simulation benchmark comprises 12 diverse manipulation tasks spanning 5 distinct environments. To rigorously assess generalization capabilities, we extend our evaluation to 3 auxiliary tasks featuring varied object geometries and environmental conditions, complemented by comprehensive ablation studies that dissect each component’s contribution. Furthermore, to validate real-world applicability and handle more complex manipulation scenarios, we deploy ViTaS within IL paradigm for challenging real-world tasks, thereby establishing its practical viability beyond simulation environments. Our experimental results demonstrate that ViTaS achieves state-of-the-art performance compared with existing visuo-tactile learning baselines.

In summary, our contributions are as follows:

- We improve the traditional contrastive learning method and use it for the fusion of visual and tactile modalities.
- We propose ViTaS, a simple yet effective representation learning paradigm that can integrate visual and tactile inputs through soft fusion contrastive as well as CVAE, utilizing it to guide the training of visuomotor learning.
- We evaluate our algorithm on various tasks in both simulation and real-world environment, demonstrating state-of-the-art performance against various baselines.

II. RELATED WORK

a) Visuo-Tactile Representation Learning: In recent years, numerous cross-modal representation learning methods have emerged, particularly those focused on visuo-tactile

integration, as demonstrated by [5], [6], [12]–[20]. Among them, [21] utilizes an adversarial loss to learn representation in the latent space, while [9] leverages a transformer architecture to integrate multiple modalities, introducing alignment and contact loss to enhance performance. [8] proposes a jointly visuo-tactile training scheme using an MAE-based encoder trained through a reconstruction process.

Despite the success of these approaches in specific tasks, they often fail to fully exploit the correspondence between visual and tactile modalities, leading to suboptimal benchmark performance. In contrast, our method employs a simpler yet highly effective CNN-based encoder to improve the alignment and fusion of modalities, achieving superior performance across multiple tasks.

b) Contrastive Learning: Extended into computer vision by the MoCo series [22], [23] and SimCLR [24], contrastive learning has emerged as a prominent technique for representation learning. We intend to extend the contrastive learning paradigm to a visuo-tactile framework for reinforcement learning. Related examples include [25], [14], [26], [27], [28], [10], [15], [25], [29] and [30]. Among the works most closely related to ours, [14] proposes a visuo-tactile fusion approach based on contrastive pre-training. [30] employs contrastive loss within the visual modality to enhance policy learning. [15] incorporates tactile, vision, and text using contrastive learning to solve downstream tasks.

However, as [10] mentions, simply doing instance discrimination tends to neglect some key information since two resembling samples may be negatives for each other due to distinct timesteps. The phenomenon also pops up in the field of cross-modal contrastive learning. We refine the contrastive learning method to alleviate the issue and better integrate different modalities, which is elaborated in Section III-A.

III. METHOD

In this section, we elaborate **Visual Tactile Soft Fusion Contrastive Learning** (ViTaS), an advanced visuo-tactile fusion framework tailored for visuomotor learning illustrated

in Figure 2. To exploit the similarity between visual and tactile information, we propose *soft fusion contrastive learning* (Section III-A), an extension of the conventional contrastive learning paradigm. Furthermore, to leverage their complementarity and handle the self-occluded scenarios, we integrate both modalities to guide the CVAE (Section III-B). Next, we elaborate these main designs with details.

A. Soft Fusion Contrastive Learning

We define a trajectory as $\Gamma = \{o_i, t_i\}_{i=1}^N$, where o_i and t_i denote the visual and tactile observations at timestep i , respectively. The encoders $f_o(\cdot)$ and $f_t(\cdot)$ map these observations into their latent representations. For simplicity, we refer to o_i and t_i at the same timestep as *dual* counterparts.

Inspired by [10], we introduce *soft fusion contrastive learning*, a novel cross-modal contrastive paradigm designed to enhance multi-modal fusion. For brevity, we refer to it as soft fusion contrastive. The central intuition is that visual and tactile representations often exhibit strong structural alignment in non-occluded scenarios, making them particularly well-suited for contrastive learning.

In the first stage, given an image observation o_i , we retrieve its top- K most similar images o_{i_1}, \dots, o_{i_K} . The corresponding *dual* tactile embeddings t_{i_1}, \dots, t_{i_K} are then designated as positive samples for t_i , while all other tactile maps are treated as negatives. With these positive and negative pairs, we apply a contrastive learning objective to train the encoder. Crucially, during this stage, only the tactile encoder is updated, while the visual encoder remains frozen. The resulting contrastive loss $\mathcal{L}_{\text{CON},1,i}$ is formulated in Eq. (1).

$$\begin{cases} \mathcal{L}_{\text{CON},1,i} = -\log \frac{\mathcal{F}_{1,\mathcal{P}_{1i}}(i)}{\mathcal{F}_{1,\mathcal{P}_{1i}}(i) + \mathcal{F}_{1,\mathcal{N}_{1i}}(i)} \\ \mathcal{F}_{1,S}(i) = \sum_{k \in S} \exp(f_t(t_k) \cdot f_t(t_i) / \tau) \end{cases} \quad (1)$$

Formally, \mathcal{P}_{1i} denotes the positive set of t_i , defined as the top- K most similar visual samples formulated in Eq. (2), while $\mathcal{N}_{1i} = S \setminus \mathcal{P}_{1i}$ denotes the negatives, where S is the full sample set. The operator $\text{topKmax}_k(f_o(o_i) \oplus f_o(o_k))$ retrieves the top- K most similar elements to the visual feature of o_i . The similarity between x and y is measured via cosine similarity, denoted as $x \oplus y$. K is a hyperparameter set to 10, with its ablation results presented in Section IV-D.

$$\mathcal{P}_{1i} = \{j | (f_o(o_j) \oplus f_o(o_i)) \in \text{topKmax}_k(f_o(o_i) \oplus f_o(o_k))\} \quad (2)$$

In the next stage, to avoid modality imbalance, we periodically swap the roles of o_i and t_i following [17], so that the visual encoder could also be updated. The corresponding objective and positive set are denoted as $\mathcal{L}_{\text{CON},2,i}$ and \mathcal{P}_{2i} , formulated in Eq. (3) and Eq. (4). We also have $\mathcal{N}_{2i} = S \setminus \mathcal{P}_{2i}$.

$$\begin{cases} \mathcal{L}_{\text{CON},2,i} = -\log \frac{\mathcal{F}_{2,\mathcal{P}_{2i}}(i)}{\mathcal{F}_{2,\mathcal{P}_{2i}}(i) + \mathcal{F}_{2,\mathcal{N}_{2i}}(i)} \\ \mathcal{F}_{2,S}(i) = \sum_{k \in S} \exp(f_o(o_k) \cdot f_o(o_i) / \tau) \end{cases} \quad (3)$$

$$\mathcal{P}_{2i} = \{j | (f_t(t_j) \oplus f_t(t_i)) \in \text{topKmax}_k(f_t(t_i) \oplus f_t(t_k))\} \quad (4)$$

To further stabilize training, we alternate between $\mathcal{L}_{\text{CON},1,i}$ and $\mathcal{L}_{\text{CON},2,i}$ using a switching schedule with period T_{switch} . Specifically, we define a coefficient sequence $u_i = 1/2 \times (1 - (-1)^{\lceil i/T \rceil}) = [1, 1, \dots, 1, 0, 0, \dots, 0, 1, 1, \dots]$, yielding a binary mask that cycles between consecutive blocks of ones and zeros. Full procedure is illustrated in the aqua block of Figure 2. The overall contrastive objective is therefore:

$$\mathcal{L}_{\text{CON}} = \sum_{i=1}^N (u_i \cdot \mathcal{L}_{\text{CON},1,i} + (1 - u_i) \cdot \mathcal{L}_{\text{CON},2,i}) \quad (5)$$

A natural counterargument is that soft fusion contrastive may be unnecessary, since temporally adjacent observations are already similar and could be used as positives. However, our ablation study (Section IV-D) demonstrates that this temporal heuristic is suboptimal compared to our approach, and we further provide an in-depth analysis of the reasons.

B. Conditional VAE Visuo-Tactile Feature Integration

VAE-based methods are widely adopted for visuo-tactile integration [31], [32]. In self-occluded scenarios, the agent should exploit the complementarity between visual and tactile modalities rather than relying solely on their alignment. Given the complementary information tactile map and image offer, our objective is to fully exploit the fusion of image and tactile embeddings, which could show superior performance to using the embeddings solely. To this end, inspired by [33], we adopt CVAE to reconstruct observations from the fused embeddings, thereby enforcing cross-modal consistency. A comprehensive depiction is presented in the pink block of Figure 2.

Specifically, we establish *condition* on the concatenated visuo-tactile feature c to reconstruct the current image frame o_{cur} . CVAE consists of an encoder $p_\theta(\cdot)$, decoder $q_\psi(\cdot)$, and visuo-tactile embedding projector $f_\phi(\cdot)$, which are parameterized by θ , ψ and ϕ separately. We use z to represent the latent variables, and the reconstructed frame \hat{o}_{cur} conditioned on visuo-tactile feature c can be expressed as:

$$\hat{o}_{\text{cur}} = q_\psi(p_\theta(o_{\text{cur}}, f_\phi(c)), f_\phi(c)) \quad (6)$$

In accordance with CVAE constraints, the target can be formulated as:

$$\mathcal{L}_{\text{VAE}} = \mathbb{E} [\|o_{\text{cur}} - \hat{o}_{\text{cur}}\|^2] + D_{\text{KL}}(p_\theta(z|o_{\text{cur}}, c) \| \mathcal{N}(0, 1)) \quad (7)$$

During training, the CVAE encoder, decoder, projector, and the image and tactile encoders are jointly optimized. Only the image and tactile encoders are retained for inference.

With two representation loss, we reach the final objective:

$$\mathcal{L} = \lambda \mathcal{L}_{\text{CON}} + \mu \mathcal{L}_{\text{VAE}} + \mathcal{L}_{\text{policy}} \quad (8)$$

where λ and μ are the coefficients to balance, and extra policy loss $\mathcal{L}_{\text{policy}}$ depends on the RL or IL paradigms we adopt. In RL case policy loss is \mathcal{L}_{PPG} , while \mathcal{L}_{DP} otherwise. The value of λ, μ is elaborated in Section IV-D.

IV. EXPERIMENTS

We evaluate our method on several contact-rich tasks in both simulation and real-world environments, in order to clarify the following questions:

- (i) Does ViTaS have the capability to solve complicated manipulation tasks (e.g. dexterous hand rotation)?
- (ii) How does ViTaS demonstrate generalization and robustness in tasks involving objects of various shapes, significant noise or different physical parameters?
- (iii) How does ViTaS perform in real-world settings?

All three questions will be elaborated in the following parts. Moreover, ablation and qualitative studies are also conducted for the understanding of components in ViTaS.

A. Simulation Environment Setup




Fig. 3: **Tasks.** Our method is evaluated on 12 simulation tasks and 3 real-world tasks, with various embodiment types.

1) *Tasks:* We conduct experiments using 12 simulated tasks, categorized into 5 primary parts shown in Figure 3 (a) to (f): (a). shadow dexterous hand tasks based on Gymnasium [34] (Pen Rotate, Block Rotate, and Egg Rotate); (b)(f). Robosuite [35]-based tasks; (c). Insertion tasks originated by [8] simulated in MuJoCo; (d). Mobile-Catch environment implemented by [36] and (e). Block Spinning task created by [29]. Beyond these foundational experiments, we introduce a series of auxiliary tasks involving altering object shapes in Lift or modifying physical parameters in Pen Rotation. The outcomes of (a)-(d), (f) environments are quantified in terms of success rate, and (e) is assessed based on training reward. (a)-(e) are trained with PPO, an effective reinforcement learning paradigm, while (f) with imitation learning is shown in Section IV-A.4.

2) *Tactile sensors:* It is crucial to integrate tactile sensors to obtain tactile data for ViTaS framework. For the 3 in-hand rotation and Wiping tasks, we employ the built-in tactile modules. For Lift, Insertion, Door Opening, Box Stack and Assembly, we employ a parallel gripper equipped with a $32 \times 32 \times 3$ tactile sensor at the contact surfaces between the gripper and the object. Among the 3 channels of the tactile map, channel 1 and 2 represent the normal force and channel 3 denotes shear force, following [37] and [6]. In the catch

and block spin tasks, we enhanced the Allegro hand and Leap hand with tactile sensors by integrating four $3 \times 3 \times 3$ sensors on each finger (located at the proximal, middle, distal, and tip segments) and one $3 \times 3 \times 3$ sensor on the palm. Sensors are zero-padded to form a $32 \times 32 \times 3$ input following [8].

3) *Comparison methods:* We compare ViTaS against 6 visuo-tactile representation learning baselines:

- M3L [8]: A visuo-tactile fusion training algorithm utilizing the MAE encoder for PPO policy learning.
- VTT [9]: A visuo-tactile fusion training method rooted in the transformer architecture.
- PoE [12]: A VAE-like framework to fuse two modality.
- Concatenation [13]: A multi-modal fusion method with contrastive method used to help training.
- MViTAC [14]: A framework based on contrastive pre-training to fuse two modality, abbreviated as MVT.
- ConViTaC [38]: A visuo-tactile fusion method with contrastive learning, abbreviated as CVT.




Fig. 4: **ViTaS with Imitation Learning.** The imitation learning paradigm is adopted to further test feature extraction ability of ViTaS in different settings, which we use for both simulation and real world tasks.

4) *ViTaS with Imitation Learning:* To further assess ViTaS's capability in extracting and integrating multimodal features across diverse scenarios, we adopt Diffusion Policy (DP) [39], a pioneering generative framework for robotic manipulation that formulates action prediction as a conditional denoising process, for both simulation and real-world evaluations. As illustrated in Figure 4, we replace DP's original CNN and transformer encoders with those from ViTaS, while retaining the same training procedure for soft fusion contrastive and CVAE modules as in simulation. For example, within the CVAE module, images are reconstructed from combined visuo-tactile embeddings. We evaluate on 3 benchmark tasks, Box Stack, Wiping and Assembly. To ensure a fair comparison, DP receives inputs from 3 cameras, whereas ViTaS relies solely on a single head camera and tactile sensors. This setting highlights that properly leveraging tactile features can compensate for reduced image, exhibiting robustness under limited visual input.

B. Simulation Experiment Results

1) *RL and IL results*: Our algorithm is compared against 4 baseline methods across 9 primitive tasks ((a)-(e) in Figure 3) for reinforcement learning. We evaluate each algorithm in each environment 5 times under different random seeds, and average the results when training 3×10^6 timesteps to obtain the performance metrics. As for imitation learning, we train all methods across 3 harder tasks ((f) in Figure 3) for 10^3 epochs, with 50 collected expert demonstrations per task.

TABLE I: **Performance.** Each experiment repeats 5 times. Green for optimal results while purple for suboptimal.

Tasks / Methods	ViTaS	MVT	CVT	M3L	PoE	VTT	Concat
Insertion	98.2	64.3	83.9	72.1	11.4	78.6	19.3
Door	100.0	100.0	100.0	100.0	98.2	99.8	100.0
Lift	97.5	97.9	70.2	20.6	71.9	70.4	76.7
Pen Rotate	99.2	77.1	79.6	73.1	0.3	0.7	2.9
Dual Arm Lift	100.0	87.5	90.8	88.2	92.6	77.1	76.8
Mobile Catch	64.9	44.2	64.2	15.8	0.4	53.3	0.6
Egg Rotate	85.7	71.6	58.1	4.2	0.9	0.5	0.7
Block Rotate	93.3	69.1	70.4	11.6	0.8	1.3	4.4
Block Spin	70.5	35.7	43.4	30.8	20.6	0.9	15.7
Insertion Noisy	89.2	51.6	70.5	47.3	20.7	63.4	26.9
Lift w/ Cap	99.6	67.8	49.9	54.2	58.1	54.7	87.5
Lift w/ Can	97.8	76.7	55.9	41.3	52.6	69.4	75.8
Average	91.4	70.3	71.5	46.6	35.7	47.5	40.6

TABLE II: **Simulation imitation learning results.** Each experiment repeats 3 times with different random seeds.

Methods / Tasks	Box Stack	Wiping	Assembly	Average
DP w/ ViTaS	53.3	71.7	56.3	60.4
DP w/ CNN	29.7	57.0	28.0	38.2
DP w/ Transformer	30.3	45.0	25.0	33.4

TABLE III: **Generalization ability.**

Tasks / Methods	ViTaS	MVT	CVT	M3L
Pen Rotate w/ Fixed Target	99.2	77.1	79.6	73.1
Pen Rotate w/ Random Target	78.4	47.5	51.3	42.7

As the results shown in Table I, several baselines show excellent performance in simple tasks like Door Open and Insertion. However, for tough tasks like Egg Rotate and Block Rotate, which require methods to incorporate visual and tactile information jointly, few baselines can solve it within a limited horizon, while ViTaS maintains its performance. Moreover, as shown in Table II, ViTaS shows superior performance against DP in all 3 tasks, showing the huge difference touch could make in manipulation tasks. **This underscores its exceptional capability to extract features and solve complex tasks, clarifying question (i).**

2) *Generalization and robustness results*: In order to assess the generalization capability of our approach, we introduce auxiliary variants of the previously mentioned Lift and Pen Rotate tasks. the object shape in Lift task is modified from a cube to cylinder and capsule in both training and testing phases, allowing us to evaluate the method's

resilience to changes in object geometry. As for the Pen Rotate task, we randomize the target angle within a large range to test the generalization of methods. Furthermore, to test the robustness, we also add Gaussian noise with 0.3 noise level in Insertion task. The intensity of Gaussian noise of different noise levels could refer to Figure 6. All other settings are in alignment with the preceding 9 tasks.

As illustrated in Table I, when the object shape is changed, every baseline model experiences a performance drop when training 3×10^6 timesteps, indicating sensitivity to these alterations. ViTaS, however, exhibits a negligible decrease, demonstrating its resilience to variations in object geometry. Moreover, as the results shown in Table III, ViTaS demonstrates superior generalization ability to other baselines, given the performance in the randomized target position of Pen Rotate task. **The results across all auxiliary tasks provide a clear answer to question (ii).**

C. Real-World Experiment




Fig. 5: **Real-World Robot Setting.**

1) *Tasks*: To better understand the overall performance of ViTaS, we develop 3 real-world experiments to show the effectiveness, shown in Figure 3 (g): (1). Dual Arm Clean (DAC). The robot has to sweep a small amount of rubbish (e.g. a piece of paper ball) to the trash can. (2). Table Pick Place (TPP). The robot has to move the bottles or cans to the coaster. This task has two settings: one (TPP-1) uses a single type of bottle and the other (TPP-3) uses three types. (3). Fridge Pick Place (FPP). The robot has to move the bottles from third level to the second of refrigerator.

2) *Experiment setup*: The overall working space is shown in Fig 5. We use Galaxea-R1 Humanoid Robot for manipulation, with tactile sensors attached to the end effector.

- Camera: we use multiple cameras to obtain RGB visual information. **Only** the head camera of Galaxea-R1 (Zed 2) is used in all ViTaS experiments, while 2 wrist cameras (RealSense D435i) are needed for better performance in some *baseline* experiments.
- Tactile sensors: We use sensors mentioned in [40] producing real-time $16 \times 16 \times 1$ discrete 1D haptic maps for tactile information. The tactile sensors are attached to the gripper, obtaining tactile maps during data collecting and inference. A snapshot shown in Figure 5 demonstrates that when grippers contact objects, corresponding position of tactile maps have numerical changes (displayed in light colors) compared to others.
- Data collection: We collect 100 real-world expert trajectories for each task, using Meta Quest 3 to teleoperate.

3) *ViTaS in real world*: As mentioned in Section IV-A.4, we also adopt imitation learning paradigm for real-world experiments, where the pipeline is shown in Figure 4. It is noteworthy that the real-world setting differs slightly from the simulation environment, requiring minor adaptations to achieve better empirical performance. To accommodate the 1D haptic map produced by our tactile sensor, the CNN tactile encoder is adjusted accordingly from $32 \times 32 \times 3$ to $16 \times 16 \times 1$. Both ViTaS and DP are trained with real data for 10^3 epochs, and are then used for real-world R1 humanoid to calculate success rate in manipulation tasks. Since we use real data for training, there is no sim-to-real issues for ViTaS.

4) *Comparing method*: In resemblance with the simulation experiment settings, we compare ViTaS with DP in the real world. Given the CNN encoder shows better results in Table II, we use learning-from-scratch CNN encoder for DP to assess the performance. Notably, ViTaS uses only the *head* camera and tactile sensor as input, whereas DP leverages all three cameras (*head, left, right*), thus receiving richer visual information. This comparison highlights that ViTaS achieves superior performance despite having access to less sensory input, further demonstrating its effectiveness.

5) *Results*: We have done Table Pick Place and Fridge Pick Place with 25 repetitions, with the target position shifted slightly each time. We also complete Dual Arm Clean with 10 repetitions, placing only *one* piece of litter during benchmarking. The results are shown in table IV.

TABLE IV: **Real-world experiment results.**

Method / Tasks	DAC	TPP-1	TPP-3	FPP	Average
ViTaS	30.0	42.0	36.0	76.0	46.0
DP	20.0	36.0	24.0	40.0	30.0

ViTaS consistently outperforms DP across 3 real-world tasks, achieving an average success rate improvement of 16%, despite relying on reduced camera input. Additionally, a comparative analysis of tasks TPP-1 and TPP-3 demonstrates ViTaS’s superior generalization, with a smaller performance drop between tasks compared to DP. Notably, occlusions induced by the robotic arm in the head camera’s field of view can impair visual perception. Nevertheless, ViTaS maintains robust performance, suggesting that tactile sensing effectively compensates for these visual limitations. This resilience highlights the efficacy of the CVAE mechanism in handling self-occlusion scenarios, underscoring ViTaS’s ability to integrate both modalities seamlessly. Given the better performance in real-world settings, **we reach the answer to question (iii).**

D. Ablation Study

To verify the fidelity of each component in ViTaS, we conduct extensive ablation experiments, showing the necessity of different designs. The general ablation results are presented in Table V, where we use abbreviations of experiments in the first row, corresponding to ViTaS, w/o. **T**Actile information, w/ **U**nified Encoder, w/o. **S**oft fusion contrastive module, w/o. **C**VAE, w/ **T**ime Contrastive, $K = 1$, $K = 20$ and

$K = 50$. The results are shown in columns V, TA, U, S, C, TC, K1, K20, K50, respectively. Detailed analysis of each experiment is clarified in the following sections.

TABLE V: **Ablation study.** Each experiment repeats 5 times. Green for optimal results while purple for suboptimal.

Tasks	V	TA	U	S	C	TC	K1	K20	K50
Insertion	99.2	88.1	61.6	90.3	71.1	75.2	83.3	85.1	78.7
Block Rotate	92.7	67.7	18.4	67.7	70.2	79.5	88.0	77.9	70.1
Egg Rotate	85.3	24.3	3.3	6.5	48.4	57.7	65.2	41.2	3.6
Average	92.5	60.9	27.1	54.7	63.2	70.6	78.8	68.1	67.3

Is tactile information crucial? We conduct 2 main experiments in this part. First we eliminate the tactile information, retaining only the visual data, and solely utilize the image encoder, while handling the corresponding tactile information through zero-padding. Additionally, the workflow outlined in [8] employs a unified MAE encoder across both modalities which overlooks their inherent differences and may result in less discriminative features and reduced effectiveness. To further demonstrate that tactile maps provide complementary information requiring separate encoders, we also design an experiment where visual and tactile inputs are directly concatenated and processed by a shared encoder.

The TA-column in Table V show that when ablating tactile information, the success rate drops 32% on average. Thus, we prove that tactile information is crucial in manipulation tasks. Using a unified encoder, however, is not a good choice either, given the poor performance in the U-column in Table V, especially for the 2 rotation tasks. We attribute this result to the fact that unified encoder integrate both modalities roughly, neglecting the inherent discrepancy of visual and tactile maps which separate encoders could alleviate.

How much do CVAE and soft fusion contrastive contribute to ViTaS? In order to clarify the effectiveness of each component, we remove the CVAE and soft fusion contrastive separately, conducting experiments on the same benchmarks. Results of ViTaS without CVAE or soft fusion contrastive are shown in S and C-column of Table V. The performance drops heavily without CVAE or soft fusion contrastive learning, showing the necessity of two main designs.

K in soft fusion contrastive learning. We explore the impact of varying K , for instance, setting it to 1, 10 (ours), 20 and 50, to observe how the results are affected. It is noteworthy that image and tactile at the same timestep are the only positives for each other when $K = 1$, adopting the same process as conventional cross-modal contrastive learning. Therefore, by comparing results between ours and $K = 1$, we can also clarify whether soft fusion contrastive could outperform conventional contrastive learning method.

The last 3 columns of Table V show the effectiveness of different K in ViTaS. The results when $K = 1$ show that though conventional contrastive learning can achieve relatively excellent performance, it still has gap with ViTaS (i.e. $K = 10$), while too large K also causes performance drops.

Soft fusion contrastive vs. time contrastive. To verify

the effectiveness of soft contrastive in another perspective, we carry out experiments utilizing an alternative contrastive approach, namely *time contrastive*, to highlight the indispensable role of cross-modal soft fusion contrastive learning. Neighboring frames (i.e. a fixed number of preceding and succeeding frames) are treated as positives in this method, while distant frames serve as negatives, echoing with [41]. The motivation behind the ablation lies in that, as mentioned in Section III-A, despite frames within close time intervals often appearing to be similar, it is crucial during the contrastive learning process to identify the K most analogous frames, which may not necessarily be temporally adjacent. This distinction underscores the importance of going beyond mere time contrastive. As shown in TC-column in table V, soft fusion contrastive learning is necessary as it outperforms time contrastive.

Weight ablation in learning objective. In Eq. (8), λ, μ indicate the weights of soft fusion contrastive and CVAE modules, and we initially set $\lambda = 1, \mu = 0.1$. To confirm that these parameters are indeed optimal, we conduct an ablation study by systematically varying the coefficients and evaluating the performance of our method on two particularly challenging tasks. As shown in Table VI, the performance when $\lambda = 1, \mu = 0.1$ is superior to other settings, showing the rationale of our chosen weights.

TABLE VI: **Ablation on λ, μ in learning objective.**

μ ($\lambda = 1$)	0.1 (ViTaS)	0.01	1	10	0.5	0.05
Egg Rotate	85.7	42.5	67.7	58.4	68.6	77.8
Block Rotate	93.3	60.3	74.1	48.9	80.2	76.5
λ ($\mu = 0.1$)	1 (ViTaS)	0.1	10	100	5	0.5
Egg Rotate	85.7	57.9	55.2	41.3	69.1	78.8
Block Rotate	93.3	74.0	66.9	43.2	79.8	85.1

In conclusion, our ablation study delves deep into our algorithm to analyze the effectiveness of each component. The results prove that tactile information, soft fusion contrastive learning and CVAE are of high importance, while soft fusion contrastive performs better than other contrastive methods like conventional contrastive learning and time contrastive.

E. Qualitative Analysis

To demonstrate the impact of the ViTaS, we use weights in the Egg Rotate task to reconstruct images from pure Gaussian noise conditioned on visuo-tactile embeddings. We compare performance under different noise levels in the observation space (visual and tactile) against the token-based MAE in M3L. Moreover, to validate CVAE’s role in cross-modal complementarity, we ablate the two main designs respectively by blocking gradients of soft fusion contrastive or CVAE propagating back to the image and tactile encoders, while keeping the CVAE encoder-decoder trainable. We then reconstruct occluded images given concatenated visuo-tactile features as condition from Gaussian noise.

As illustrated in Figure 6 (a), the results indicate that ViTaS surpasses the token-based MAE in reconstructing critical interaction details, such as the egg’s location, which




Fig. 6: **Reconstruction visualization.** (a) compares the reconstruction quality of ViTaS and M3L under varying noise levels. (b) presents ViTaS reconstructions under heavy noise applied to different modalities. (c) illustrates reconstruction performance under masking when ablating key designs.

is vital for the task. Our method also maintains robust under higher level of noise, underscoring the high quality of the visuo-tactile embeddings used as conditions. In Figure 6 (b), applying heavy noise (noise level 0.5) to one modality yields better reconstructions than adding noise to both, demonstrating vision-touch complementarity. Figure 6 (c) shows that even with core image regions masked, ViTaS reliably reconstructs observations when guided by tactile and masked visual features, validating CVAE’s role in handling self-occlusion and complementarity between two modalities.

V. CONCLUSION AND LIMITATIONS

In general, we present ViTaS, a concise yet effective visuo-tactile fusion framework. Analogous to human physiology, it extends visual and tactile perception to RL and IL, achieving strong results in both simulation and real-world experiments. Specifically, *soft fusion contrastive learning* extracts key features across modalities, while a CVAE module exploits their complementarity. Real-world experiments confirm ViTaS’s effectiveness, and ablation study plus qualitative analysis highlight the necessity of each component.

Despite its success, ViTaS faces 2 main limitations. The first is that due to the physical capability and the bottleneck of RL in high-dimensional cases, some high dynamic accurate manipulation like pen spinning in real world remains challenging. The other is that the potential of visuo-tactile sensing in deformable object manipulation [42] warrants extra exploration. In the future, we will further explore the ability of fused visuo-tactile features in more complex scenarios, with advanced simulation and real-world platforms.

REFERENCES

- [1] P. Apkarian-Stielau and J. M. Loomis, “A comparison of tactile and blurred visual form perception,” *Perception & Psychophysics*, vol. 18, pp. 362–368, 1975.

[2] A. M. Kappers, "Human perception of shape from touch," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 366, no. 1581, pp. 3106–3114, 2011.

[3] Z. Yuan, G. Ma, Y. Mu, B. Xia, B. Yuan, X. Wang, P. Luo, and H. Xu, "Don't touch what matters: Task-aware lipschitz data augmentation for visual reinforcement learning," *arXiv preprint arXiv:2202.09982*, 2022.

[4] J. Pitz, L. Röstel, L. Sievers, and B. Bäuml, "Dextrous tactile in-hand manipulation using a modular reinforcement learning architecture," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 1852–1858.

[5] T. Lin, Y. Zhang, Q. Li, H. Qi, B. Yi, S. Levine, and J. Malik, "Learning visuotactile skills with two multifingered hands," *arXiv preprint arXiv:2404.16823*, 2024.

[6] J. Xu, S. Kim, T. Chen, A. R. Garcia, P. Agrawal, W. Matusik, and S. Sueda, "Efficient tactile simulation with differentiability for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2023, pp. 1488–1498.

[7] H. Qi, A. Kumar, R. Calandra, Y. Ma, and J. Malik, "In-hand object rotation via rapid motor adaptation," in *Conference on Robot Learning*. PMLR, 2023, pp. 1722–1732.

[8] C. Sferrazza, Y. Seo, H. Liu, Y. Lee, and P. Abbeel, "The power of the senses: Generalizable manipulation from vision and touch through masked multimodal learning," 2023.

[9] Y. Chen, M. Van der Merwe, A. Sipos, and N. Fazeli, "Visuo-tactile transformers for manipulation," in *6th Annual Conference on Robot Learning*, 2022.

[10] T. Han, W. Xie, and A. Zisserman, "Self-supervised co-training for video representation learning," *Advances in neural information processing systems*, vol. 33, pp. 5679–5690, 2020.

[11] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," *Advances in neural information processing systems*, vol. 28, 2015.

[12] M. A. Lee, Y. Zhu, P. Zachares, M. Tan, K. Srinivasan, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Learning multimodal representations for contact-rich tasks," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 582–596, 2020.

[13] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks," in *2019 International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 8943–8950.

[14] V. Dave, F. Lygerakis, and E. Rueckert, "Multimodal visual-tactile representation learning through self-supervised contrastive pre-training," *arXiv preprint arXiv:2401.12024*, 2024.

[15] F. Yang, C. Feng, Z. Chen, H. Park, D. Wang, Y. Dou, Z. Zeng, X. Chen, R. Gangopadhyay, A. Owens *et al.*, "Binding touch to everything: Learning unified multimodal tactile representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 26340–26353.

[16] H. Li, Y. Zhang, J. Zhu, S. Wang, M. A. Lee, H. Xu, E. Adelson, L. Fei-Fei, R. Gao, and J. Wu, "See, hear, and feel: Smart sensory fusion for robotic manipulation," *arXiv preprint arXiv:2212.03858*, 2022.

[17] 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.

[18] S. Li, Z. Wang, C. Wu, X. Li, S. Luo, B. Fang, F. Sun, X.-P. Zhang, and W. Ding, "When vision meets touch: A contemporary review for visuotactile sensors from the signal processing perspective," *arXiv preprint arXiv:2406.12226*, 2024.

[19] S. Li, X. Yin, C. Xia, L. Ye, X. Wang, and B. Liang, "Tata: A universal jamming gripper with high-quality tactile perception and its application to underwater manipulation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6151–6157.

[20] Y. Lu, Y. Tian, Z. Yuan, X. Wang, P. Hua, Z. Xue, and H. Xu, "H3 dp: Triply-hierarchical diffusion policy for visuomotor learning," *arXiv preprint arXiv:2505.07819*, 2025.

[21] Y. Li, J.-Y. Zhu, R. Tedrake, and A. Torralba, "Connecting touch and vision via cross-modal prediction," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10609–10618.

[22] 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.

[23] X. Chen, S. Xie, and K. He, "An empirical study of training self-supervised vision transformers," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9640–9649.

[24] 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.

[25] D. Wang and M. Hu, "Contrastive learning methods for deep reinforcement learning," *IEEE Access*, vol. 11, pp. 97107–97117, 2023.

[26] X. Yuan, Z. Lin, J. Kuen, J. Zhang, Y. Wang, M. Maire, A. Kale, and B. Faieta, "Multimodal contrastive training for visual representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6995–7004.

[27] M. Laskin, A. Srinivas, and P. Abbeel, "Curl: Contrastive unsupervised representations for reinforcement learning," in *International conference on machine learning*. PMLR, 2020, pp. 5639–5650.

[28] D. Wang and M. Hu, "Contrastive learning methods for deep reinforcement learning," *IEEE Access*, 2023.

[29] Z. Yuan, T. Wei, S. Cheng, G. Zhang, Y. Chen, and H. Xu, "Learning to manipulate anywhere: A visual generalizable framework for reinforcement learning," *arXiv preprint arXiv:2407.15815*, 2024.

[30] A. Zhan, R. Zhao, L. Pinto, P. Abbeel, and M. Laskin, "Learning visual robotic control efficiently with contrastive pre-training and data augmentation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 4040–4047.

[31] A. V. Nair, V. Pong, M. Dalal, S. Bahl, S. Lin, and S. Levine, "Visual reinforcement learning with imagined goals," *Advances in neural information processing systems*, vol. 31, 2018.

[32] C. Bai, P. Liu, K. Liu, L. Wang, Y. Zhao, L. Han, and Z. Wang, "Variational dynamic for self-supervised exploration in deep reinforcement learning," *IEEE Transactions on neural networks and learning systems*, vol. 34, no. 8, pp. 4776–4790, 2021.

[33] P. Bachhav, M. Todisco, and N. Evans, "Latent representation learning for artificial bandwidth extension using a conditional variational auto-encoder," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 7010–7014.

[34] M. Towers, A. Kwiatkowski, J. Terry, J. U. Balis, G. De Cola, T. Deleu, M. Goulão, A. Kallinteris, M. Krimmel, A. KG *et al.*, "Gymnasium: A standard interface for reinforcement learning environments," *arXiv preprint arXiv:2407.17032*, 2024.

[35] Y. Zhu, J. Wong, A. Mandlekar, R. Martín-Martín, A. Joshi, S. Nasiriany, and Y. Zhu, "robosuite: A modular simulation framework and benchmark for robot learning," *arXiv preprint arXiv:2009.12293*, 2020.

[36] Y. Zhang, T. Liang, Z. Chen, Y. Ze, and H. Xu, "Catch it! learning to catch in flight with mobile dexterous hands," *arXiv preprint arXiv:2409.10319*, 2024.

[37] I. H. Taylor, S. Dong, and A. Rodriguez, "Gelslim 3.0: High-resolution measurement of shape, force and slip in a compact tactile-sensing finger," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 10781–10787.

[38] Z. Wu, Y. Zhao, and S. Luo, "Convitac: Aligning visual-tactile fusion with contrastive representations," *arXiv preprint arXiv:2506.20757*, 2025.

[39] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, p. 02783649241273668, 2023.

[40] B. Huang, Y. Wang, X. Yang, Y. Luo, and Y. Li, "3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing," *arXiv preprint arXiv:2410.24091*, 2024.

[41] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, S. Levine, and G. Brain, "Time-contrastive networks: Self-supervised learning from video," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 1134–1141.

[42] C. Jiang, W. Xu, Y. Li, Z. Yu, L. Wang, X. Hu, Z. Xie, Q. Liu, B. Yang, X. Wang *et al.*, "Capturing forceful interaction with deformable objects using a deep learning-powered stretchable tactile array," *Nature Communications*, vol. 15, no. 1, p. 9513, 2024.