

Tactile Memory with Soft Robot: Robust Object Insertion via Masked Encoding and Soft Wrist

Abstract—Tactile memory, the ability to store and retrieve touch-based experience, is critical for contact-rich tasks such as key insertion under uncertainty. To replicate this capability, we introduce Tactile Memory with Soft Robot (TaMeSo-bot), a system that integrates a soft wrist with tactile retrieval-based control to enable safe and robust manipulation. The soft wrist allows safe contact exploration during data collection, while tactile memory reuses past demonstrations via retrieval for flexible adaptation to unseen scenarios. The core of this system is the Masked Tactile Trajectory Transformer (MAT³), which jointly models spatiotemporal interactions between robot actions, distributed tactile feedback, force-torque measurements, and proprioceptive signals. Through masked-token prediction, MAT³ learns rich spatiotemporal representations by inferring missing sensory information from context, autonomously extracting task-relevant features without explicit subtask segmentation. We validate our approach on peg-in-hole tasks with diverse pegs and conditions in real-robot experiments. Our extensive evaluation demonstrates that MAT³ achieves higher success rates than the baselines over all conditions and shows remarkable capability to adapt to unseen pegs and conditions.

I. INTRODUCTION

A tactile memory system, which stores and retrieves tactile information, is essential for daily activities [1]. For example, humans can insert a key even when the exact location of the lock is uncertain by exploring the surrounding area through gentle touch and relying on the tactile memory of how the key fits and inserts into the lock. This ability to handle uncertainty through tactile feedback is a fundamental aspect of human robust manipulation. While this process is trivial for humans, replicating such intuitive and adaptive behavior in robots remains a major challenge. This study explores a robotic system, **Tactile Memory with Soft Robot (TaMeSo-bot)**, that integrates two key elements: soft robots, which enable stable contact, and a tactile memory capable of storing and retrieving tactile information within a database.

The goal of this research is to achieve robust peg-in-hole insertion, a widely studied yet challenging contact-rich manipulation task [2], [3], through a tactile memory system. As depicted in Fig. 1, robustness refers to exhibiting adaptive behavior across unseen pegs and conditions. TaMeSo-bot enables flexible adaptation to novel scenarios by storing and retrieving past tactile demonstrations, using a largely deformable soft wrist [4] for both safe data collection and robust execution. Unlike prior approaches that rely on predefined contact state sequences [5], [6], our system requires no manual annotation of subtasks such as fit, align, and insert. Instead, TaMeSo-bot uses tactile cues to retrieve and execute appropriate actions from past demonstrations, implicitly

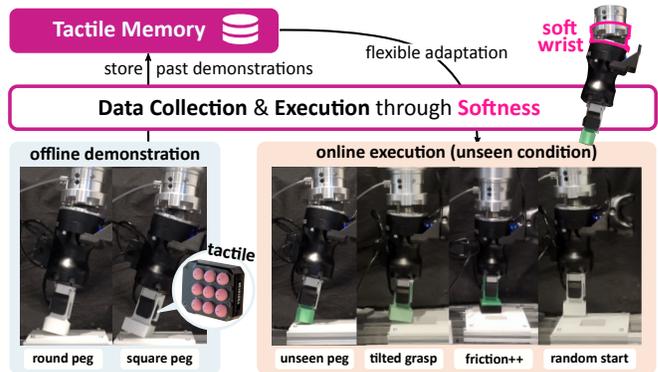


Fig. 1: **TaMeSo-bot** uses softness (i.e., mechanical compliance) to safely collect tactile demonstrations (left) and flexibly adapt to unseen conditions (right) through retrieval from tactile memory.

recognizing subtask transitions. This approach mirrors how humans unconsciously sense subtask transition through touch without explicitly recognizing each contact phase [7]. The integration of softness is also crucial: with rigid robots, slight pose errors during contact generate large contact forces that trigger protective stops, while the wrist’s softness passively accommodates misalignment, enabling safe contact exploration and stable tactile data collection. The wrist’s softness also protects the tactile sensing system by absorbing excessive contact forces during insertion. Combined with a durable distributed tactile sensor [8] that withstands the high shear forces, our system achieves reliable tactile sensing under diverse contact conditions.

While reusing prior experiences has shown substantial potential for robot manipulation [9]–[11], the challenge lies in learning rich representations. Specifically, the representation must enable effective matching of current observations to relevant past experiences based on subtle differences in sensor information, without relying on explicit state labels. Integrating tactile feedback with other modalities such as force-torque (F/T) and proprioceptive signals is helpful for capturing rich contextual information to achieve robust manipulation. However, existing methods have not directly addressed the complex spatiotemporal correlations between tactile responses and other modalities, including robot actions. They typically process modalities separately, rely on simple concatenation [11], or ignore temporal correlations across modalities [10], [12], failing to capture these spatiotemporal dependencies critical for robust contact manipulation.

To address this challenge, we design the **Masked Tactile**

Trajectory Transformer (**MAT³**) that encodes spatiotemporal interaction between tactile signals and other modalities for effective storage and retrieval in tactile memory. Fig. 2 illustrates an overview of our TaMeSo-bot equipped with **MAT³** as its core component. Inspired by neurophysiological findings that “*tactile memory encodes spatiotemporal information and interacts with other sensory modalities*” [1], we employ a distributed tactile sensor integrated into the gripper and a built-in F/T sensor on the robot. **MAT³** takes sequences of tokens derived from each taxel, action, and other modalities, including force/torque signals, soft wrist poses, and end-effector poses as inputs. Through spatial and temporal encoding, the bidirectional Transformer computes cross-modal attention across different taxel locations and time horizons, enabling the extraction of spatiotemporal features that capture the intricate relationships between distributed tactile feedback and multimodal sensory signals.

To further obtain a rich feature representation relevant to the task, we train our model using masked token prediction techniques [13], [14]. The masking strategy forces the model to reconstruct masked sensor values and actions from neighboring spatiotemporal context, compelling the model to focus on the critical contact features. By learning from incomplete input sequences, **MAT³** develops rich internal representations that capture bidirectional context dependencies for robust peg-in-hole tasks. These learned representations are stored in the tactile memory and subsequently used via retrieve-and-replay execution via softness.

We validate our approach through real-world robot experiments on peg-in-hole tasks with seven different peg shapes and various perturbation conditions, including pose uncertainty and increased friction. Within the TaMeSo-bot framework, we evaluate **MAT³** against other tactile encoders and demonstrate that **MAT³** achieves higher success rates than baselines and handle uncertainty in grasp pose and successfully adapts to previously unseen scenarios. In summary, we make two key contributions: (1) **TaMeSo-bot**, which combines soft robotics with tactile memory for robust contact peg-in-hole manipulation and (2) **MAT³**, a masked bidirectional-transformer encoder that extracts spatiotemporal features from distributed taxels and other modalities.

II. RELATED WORK

A. Soft robots and tactile sensors for contact

Soft robots equipped with tactile sensors can safely interact with objects and infer their properties [15], and are therefore widely used for grasping and other contact-rich manipulation tasks. Soft fingers and palms with tactile sensing have demonstrated stable grasping [16]–[20]. Whisker-inspired soft sensors support tactile exploration and object localization [21]–[23]. Tendon-driven silicone grippers can estimate contact forces from tendon tension measurements and have been applied to contact-rich hair manipulation [24]. Pneumatic grippers with tactile sensing have been used to manipulate thin and deformable objects [25]. This study tackles peg-in-hole tasks that require stable contact and high-precision manipulation. While prior work has addressed insertion with

soft robots and tactile sensing, existing approaches typically rely on heuristic subtask segmentation [26] or simulation environments [27]. In contrast, we present a retrieval-based method that requires neither.

B. Tactile Representation Learning for Manipulation

Tactile sensing captures the fine-grained interaction forces between robots and their environments that underpin dexterous manipulation [28]. Tactile representation learning methods have been developed to extract essential features from spatiotemporal data obtained through vision-based or distributed tactile sensors. Spatial structure in high-resolution tactile arrays is typically encoded with CNNs [10], [11], Vision Transformer (ViT)-style Transformers [12], [29], [30], or graph convolutions over taxel grids [31]–[33]. To add *temporal* information, these spatial encoders are often paired with recurrent layers such as LSTMs that process the frame-wise embeddings in sequence [34], [35]. These approaches, however, model spatial and temporal aspects separately rather than capturing their joint spatiotemporal correlations. In this work, we propose to learn a spatiotemporal tactile representation with a bidirectional Transformer that enables (1) jointly modeling spatiotemporal correlations between distributed tactile signals and concurrent actions, (2) learning rich representations focused on critical features through masked token prediction, and (3) incorporating auxiliary modalities through weighted soft-concatenation that provides tactile-centric modality fusion.

C. Retrieval from Offline Demonstrations

A growing body of work explores retrieval-enhanced learning to control, which leverages large, unlabeled demonstration databases to retrieve relevant past experiences for training policies [36]–[38] or decision-making [39]. This paradigm has shown promising results while eliminating the need for explicit task segmentation or manual annotation. In parallel, another line of work [11], [40], [41] pursues a non-parametric control policy via retrieve-and-replay of demonstrations without learning a parametric policy. While effective in some scenarios, these methods still require learned alignment modules or residual policies to compensate for the mismatch between retrieved demonstrations and real-world execution, adding system complexity. In contrast to prior work, **MAT³** does not require learned alignment modules due to its spatiotemporal structure and masking strategy, enabling robust action retrieval.

In summary, by combining Spatiotemporal tactile representation learning with inherent physical adaptability, our control pipeline offers a lightweight and practical alternative to existing retrieval-based frameworks for peg-in-hole tasks.

III. TAMESO-BOT

This section details the core components of the tactile memory with soft robot (TaMeSo-bot).

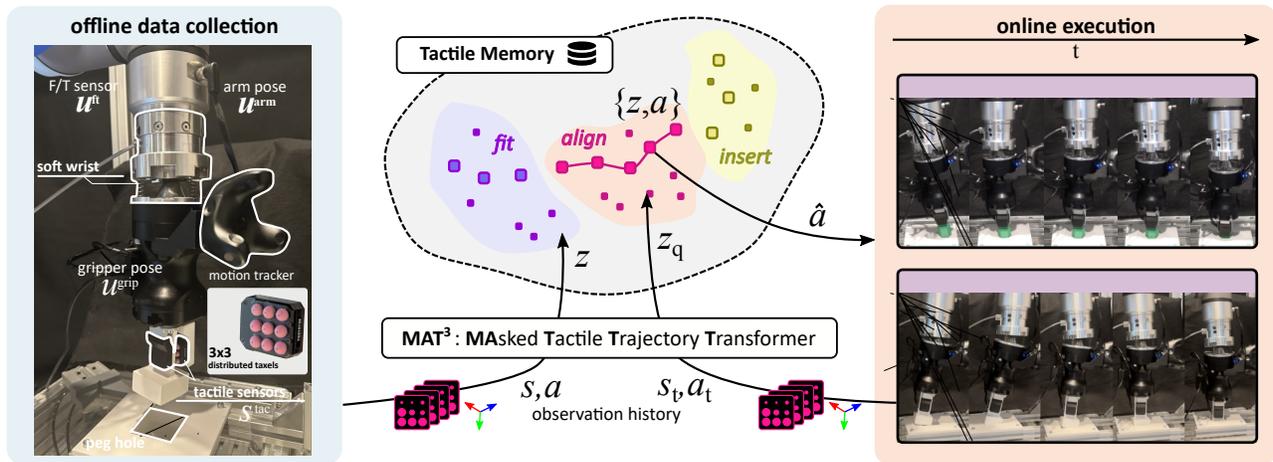


Fig. 2: Overview of the **TaMeSo-bot** system. The tactile memory system stores and retrieves tactile information from a database of encoded demonstrations, enabling robust peg-in-hole manipulation by matching current sensory inputs to similar past experiences.

A. Overview

Fig. 2 provides an overview. Our objective is to explore the synergistic potential of a tactile memory system with a soft wrist for robust contact-rich manipulation, validated through peg-in-hole tasks. We employ a rigid robot arm with an integrated F/T sensor, a soft wrist [4] and a two-finger gripper with a distributed tactile sensor (a 3×3 array of taxels) [8], each measuring a 3D force. The distributed tactile sensor supplies rich and localized contact information that is particularly valuable during peg-in-hole insertion. We assume access to the gripper pose via a motion capture system.

The soft wrist consists of three coil springs in parallel and provides a large 6D deformation, which is essential for safe data collection: only a rigid arm tends to trigger protective stops upon contact. Mechanical compliance enables safe data collection and execution by absorbing large contact forces.

We collect offline demonstration using this system via teleoperation. We then train MAT^3 , incorporating spatiotemporal representations from multimodal sensor streams with a Transformer encoder augmented by masked-token prediction (Sect. III-B), and subsequently construct and query a database for retrieval-based action selection (Sect. III-C).

B. MAT^3 : Masked Tactile Trajectory Transformer

The tactile memory system implements a non-parametric control policy that retrieves the appropriate action by querying a database whose keys are tactile representations augmented with multimodal contexts. To build these query keys, we aim to learn an encoder \mathcal{E} that maps a sub-trajectory $\tau = (\{s_{t-H+1}, a_{t-H+1}\}, \dots, \{s_t, a_t\})$ consisting of state-action pairs $\{s, a\}$ over a history window H , to a compact representation vector z_t that is later stored in a database with the corresponding action a_t . Each embedding z_t is stored in the database together with its accompanying action a_t and later used to retrieve the appropriate action at execution time.

a) *Modality Fusion and Spatio-Temporal Encoding*: At each timestep t , we collect tactile signals from a PapillArray

Tactile Sensor [8] with nine (3×3) taxels, each producing 3D force, denoted as $S_t^{\text{tac}} = \{s_t^1, \dots, s_t^9\} \in \mathbb{R}^{3 \times 9}$. As auxiliary states, we capture the F/T measurements $u_t^{\text{ft}} \in \mathbb{R}^6$, arm pose $u_t^{\text{arm}} \in \mathbb{R}^7$, and gripper pose $u_t^{\text{grip}} \in \mathbb{R}^7$ measured from a motion capture system. We also record the robot action $a_t \in \mathbb{R}^6$, which is the displacement of the 3D position $[\Delta x, \Delta y, \Delta z]$ and euler rotation representations $[\Delta \theta_x, \Delta \theta_y, \Delta \theta_z]$ of the robot arm.

Fig. 3 left summarizes how we integrate each modality and spatiotemporal information into an input embedding. The individual taxels and action serve as **base tokens** $\{s^1, \dots, s^9, a\}$, while the other sensory information $\{u^{\text{ft}}, u^{\text{arm}}, u^{\text{grip}}\}$ serves as **auxiliary tokens** that compensate for global contexts. These tokens are projected to a set of d_e -dimensional embeddings $b_t = \{b^1, \dots, b^9, b^{10}\}$ and $e_t^{\text{aux}} = \{e_t^{\text{ft}}, e_t^{\text{arm}}, e_t^{\text{grip}}\}$ by a linear projection $f(\cdot)$. We employ sinusoidal encoding for spatiotemporal modeling. Spatial relationships between taxels are captured by assigning position embeddings $y^i \in \mathbb{R}^{d_{\text{pos}}}$ to each taxel at 3×3 grid coordinates $(r_i, c_i) \in \{-1, 0, 1\}^2$ and concatenated to the respective tactile embeddings b^i . To inform temporal positions across multiple timesteps, temporal position is encoded as $e_t^{\text{time}} \in \mathbb{R}^{d_e}$. The action token is considered at the center of this grid (0,0) to facilitate spatial interaction between actions and surrounding tactile feedback. To incorporate these auxiliary embeddings e^{aux} as global context, we add their weighted sum to each base token embedding b . In summary, each input embedding $x^i \in \mathbb{R}^{d_e + d_{\text{pos}}}$ integrates information through a two-level fusion of global and local context, $x_t^i = [b_t^i + e_t^{\text{aux}}] \oplus y^i$, where

$$e_t^{\text{aux}} = w^{\text{ft}} e_t^{\text{ft}} + w^{\text{arm}} e_t^{\text{arm}} + w^{\text{grip}} e_t^{\text{grip}} + w^{\text{time}} e_t^{\text{time}}. \quad (1)$$

\oplus denotes hard-concatenation of positional embeddings and $+$ denotes soft-concatenation with weights w for auxiliary modalities and temporal information.

b) *Bidirectional Transformer-based Tactile Encoder*: We introduce a Bidirectional Transformer-based tactile rep-

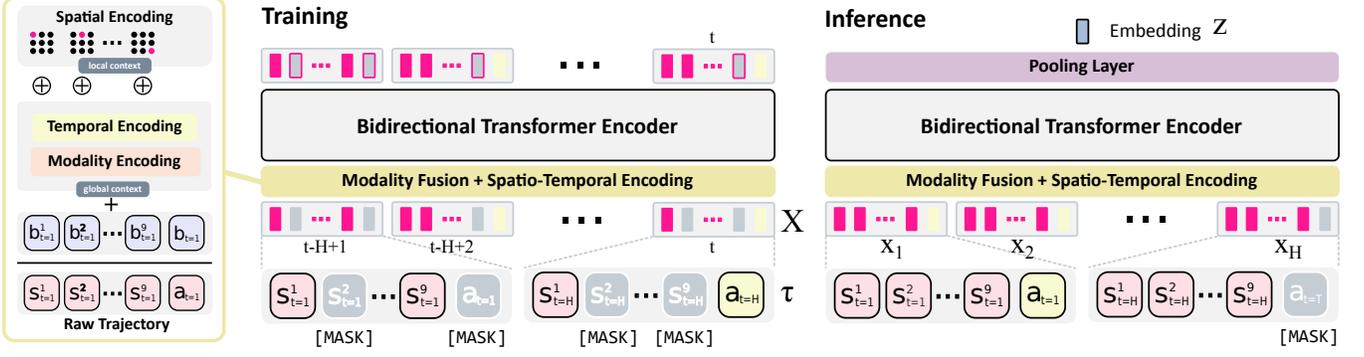


Fig. 3: **Masked Tactile Trajectory Transformer (MAT³)** for distributed taxels. We integrate auxiliary and spatiotemporal information as soft and hard-concatenation of the embeddings. During training, the encoder learns to reconstruct the states and actions within a time window H , while randomly masking input tokens. After training, the tactile trajectory datasets are encoded into a single representative embedding \mathbf{z} that captures the spatio-temporal dynamics of the tactile-action sequences.

representation learning approach that captures spatiotemporal patterns across sensory inputs and actions. Our model processes sub-trajectory sequences τ from historical observations. Through the embedding process defined in Eq. (1), a sub-trajectory τ is transformed into an input tensor $\mathbf{X} = (\{\mathbf{x}_1^1, \dots, \mathbf{x}_1^{10}\}, \dots, \{\mathbf{x}_H^1, \dots, \mathbf{x}_H^{10}\}) \in \mathbb{R}^{10 \times H \times d}$, where H is the sequence length, 10 represents the number of tokens per timestep consisting of nine taxel tokens and one action token, and $d = d_e + d_{\text{pos}}$ is the embedding dimension. We employ a bidirectional transformer encoder \mathcal{E} with L layers to process the input tensor \mathbf{X} . Each layer applies multi-head self-attention followed by a feed-forward network, enabling the model to learn complex relationships between taxels, actions, and auxiliary modalities across the temporal time window H . This architecture effectively captures the spatiotemporal structure of tactile information while maintaining computational efficiency through the attention mechanism’s ability to model long-range dependencies. Details of the model parameters are provided in the Appendix.

c) Masked Tactile Trajectory Encoder: To learn tactile representations that capture the intrinsic structure of the data, we employ masked token prediction strategy [13], [14]. As illustrated in Fig. 3, we train a Bidirectional Transformer-based encoder that learns to reconstruct input states and actions with a temporal window H while randomly masking input tokens. Each modality is first passed through a dedicated linear layer that projects its raw shape to the Transformer’s common embedding dimension d . During training, a portion of the input sequence is randomly masked as $\text{Masked}(\mathbf{X})$, forcing the model to infer missing values from neighboring spatial and temporal context. This masking strategy compels the model to identify essential features by learning critical patterns for reconstruction, while developing context-sensitive representations across different modalities.

Specifically, for each input \mathbf{X} , we sample a masking ratio uniformly from the range $[0, 0.6]$, following prior work [13]. During inference, since the current action \mathbf{a}_t is unknown, we mask the action token in the current timestep and use the

encoder to generate query vectors that capture the current tactile context. These learned representations are then stored in the tactile memory database and used for action retrieval during execution.

d) Learning Objective: After encoding, each processed token is passed through a modality-specific linear layer that projects the embedding back to the original signal dimension. The reconstructed outputs are then compared with the raw observations at the same temporal indices. The model is trained to reconstruct these observations by minimizing a mean-squared-error (MSE) loss averaged over every element in the entire mini-batch:

$$\mathcal{L} = \mathcal{L}_{\text{tactile}} + \mathcal{L}_{\text{action}}, \quad (2)$$

where $\mathcal{L}_{\text{tactile}}$ and $\mathcal{L}_{\text{action}}$ denote the MSE losses for tactile and action reconstruction, respectively, each averaged over all timesteps in the mini-batch. The loss in Eq. (2) is optimized with the Adam optimizer.

e) Feature Pooling: To obtain a global tactile representation that captures information across all tokens, we apply a pooling function $g(\cdot)$. Generally, symmetric functions serve as ideal pooling operators since they are invariant to the ordering of input tokens and can extract universal features regardless of permutation [42]. For our system, we adopt the widely used average pooling [43].

C. Database and Retrieval

The tactile memory database stores pairs of tactile representations and their corresponding actions $\{\mathbf{z}_t, \mathbf{a}_t\}$, enabling the robot to retrieve appropriate behaviors during execution based on the current sensory state.

a) Database Construction: Once an encoder \mathcal{E} is pre-trained by offline demonstrations, we construct the database using the collected trajectories. Each trajectory τ consists of a sequence of tactile tokens $\mathbf{x}_t \in \mathbb{R}^{10 \times H \times d}$, where 10 denotes the number of tokens including nine tactile sensor taxels and one action token, H is the number of timesteps to encode, and d is the feature dimension. To obtain a global

tactile representation that is invariant to the input token ordering, we perform feature pooling using a symmetric function. Following prior work [13], we adopt average pooling as our pooling strategy. This operation converts the input tensor \mathbf{x}_t into a representative vector $\mathbf{z}_t \in \mathbb{R}^d$. The resulting pair $\{\mathbf{z}_t, \mathbf{a}_t\}$, where \mathbf{a}_t is the corresponding action sequence, is then stored in the database \mathcal{D} .

b) Non-parametric Control Policy: One core hypothesis is that appropriate actions can be obtained through similarity-based retrieval from past experiences, effectively transforming the tactile memory database into a non-parametric control policy $\pi : \mathbf{s} \mapsto \mathbf{a}$. At execution time, given a sequence of tactile observation $(\mathbf{s}_{t-H+1}, \dots, \mathbf{s}_t)$ and previous actions $(\mathbf{a}_{t-H+1}, \dots, \mathbf{a}_{t-1})$, we encode it using the pretrained encoder \mathcal{E} and pooling function $g(\cdot)$ to obtain the query representation $\mathbf{z}_q \in \mathbb{R}_e^d$. Here we mask the current action token as it is unknown at the time of execution. The bidirectional attention aggregates information from the remaining tokens so that \mathbf{z}_q still captures the full tactile context. We then retrieve the k most similar representations from the database using L2 distance,

$$\text{dist}(\mathbf{z}_q, \mathbf{z}_i) = \|\mathbf{z}_q - \mathbf{z}_i\|. \quad (3)$$

For real-time robot control at 50 Hz, we employ an approximate nearest neighbor search using a hierarchical navigable small world (HNSW) graph-based index [44], which offers logarithmic-time complexity while maintaining high retrieval accuracy. From the k retrieved neighbors, we uniformly sample one of them at random. The integration of masked trajectory modeling with nearest neighbor retrieval creates a robust policy that can handle partially observed contexts and adapt to novel situations by leveraging the most relevant past experiences. Furthermore, this non-parametric approach preserves the safety characteristics of demonstration data, as the robot’s actions are bounded by the behaviors present in the offline dataset.

IV. EXPERIMENTS

We conduct real robot experiments to evaluate our TaMeSo-bot system in peg-in-hole tasks. Our experiments aim to answer the following two questions: (1) Can TaMeSo-bot complete peg-in-hole tasks? and (2) Does MAT³ improve robustness and adapt to unseen pegs and conditions?

A. Setup

a) Robot: We use a robotic arm (UR5e, Universal Robots A/S, Denmark) with an integrated F/T sensor, a parallel gripper (Hand-e, Robotiq, Canada), and a soft wrist [4], which consists of three coil springs in parallel and allows 6D deformation. The spring constant of the spring is 1.448 N/mm, and the equilibrium length is 25 mm. A motion tracker (HTC VIVE Tracker, HTC Corporation, Taiwan) is mounted on the gripper to capture its pose, and a tactile sensor (PapillArray Tactile Sensor [8], Contactile Pty Ltd, Australia) is attached to one side of the gripper.

b) Pegs and Dataset Collection: In the peg-in-hole tasks, we use seven 3D-printed pegs, including two for the training set and five for the validation set, as illustrated in Fig. 4. We collect demonstrations using the teleoperation system with a VR controller (HTC VIVE Controller, HTC Corporation, Taiwan). The robot grasps each peg in a vertical orientation. Both the sampling and control frequencies are set to 50 Hz. Each demonstration starts from one of four initial poses distributed equally around the hole and terminates when the peg is fully inserted into the hole, lasting approximately 150–300 time steps (3–6 seconds). The resulting dataset comprises 64 successful demonstrations: 32 from the square peg and 32 from the cylinder peg. The tolerances of the pegs and holes are 2 mm.

c) Baselines: Since no existing representation learning methods apply directly to our setting, which includes distributed tactile sensors, arm and gripper poses, and an F/T sensor, we compare our approach against the following baselines:

Tactile Transformer We develop a Tactile Transformer by adapting the Visuo-Tactile Transformer (VTT) [12] to our modalities. Whereas VTT uses patched RGB images as input tokens, our Tactile Transformer uses taxels. At each timestep, each taxel, action and additional modality are encoded as separate tokens and fed to a Transformer encoder with cross-modal attention. Following the original VTT, we use learned 1D position embeddings added to the tokens, whereas MAT³ uses fixed sinusoidal 2D spatial embeddings concatenated with content embeddings. The final representation is obtained by average pooling over the output tokens. Unlike MAT³, which applies soft concatenation across modalities and incorporates temporal dependencies, the Tactile Transformer maintains separated tokens per modality.

MAT³ w/o Mask This variant is identical to MAT³ but trained without the masking strategy, isolating the contribution of masked modeling.

d) Unseen Conditions: We evaluate task success rates under both seen and unseen pegs shown in Fig. 4, as well as robustness to three types of unseen perturbations illustrated in Fig. 5: (1) different starting positions, (2) increased friction between the peg and fixture surface, and (3) tilted grasping pose. For randomizing the starting position, we uniformly sample spatial offsets from the nominal starting pose used during teleoperation. To increase friction, we attach rubber tape to the peg bottom. For grasping perturbations, we manually tilt the peg by 5° before task execution.

e) Parameter Settings: We implement MAT³ as 4-layer Transformer with 8 attention heads and hidden dimension 512. We use 248-dimensional token embeddings d_e with 8-dimensional spatial embedding d_{pos} , yielding 256-dimensional input embedding. Following previous work [13], we sample the masking ratio for input tokens from a uniform distribution over the interval $[0, 0.6]$. The model is trained for 30 epochs with a batch size of 128 and a learning rate 1.0×10^{-4} .

f) Hyperparameter Selection: These parameters were selected based on preliminary experiments, where we sys-



Fig. 4: Seen and unseen peg and dimension (mm)

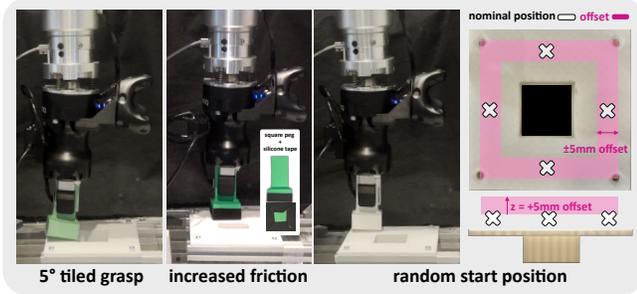


Fig. 5: Unseen conditions in the experiments.

tematically evaluated different configurations. Specifically, we tested various embedding and hidden size pairs from $\{(32, 64), (64, 128), (128, 256), (256, 512), (512, 1024)\}$, with $(256, 512)$ yielding the best action retrieval accuracy.

B. Results

Table I presents the success rates across seen and unseen pegs, with 40 trials per peg shape (10 trials from each starting position), totaling 80 trials for seen pegs and 200 trials for unseen pegs. For seen pegs, MAT^3 w/o Mask and MAT^3 achieved comparable performance at 90% and 88.8%, respectively. On unseen pegs, MAT^3 significantly outperformed MAT^3 w/o Mask with 85% versus 71% success rate, demonstrating the effectiveness of masked modeling for generalization. The Transformer baseline achieved only 22.5% and 17.5% on seen and unseen pegs, respectively, indicating that single-step cross-modal attention alone is insufficient for capturing the dynamics in this contact-rich manipulation task.

Table II evaluates robustness under three unseen perturbation conditions while MAT^3 w/o Mask and the Transformer baseline degraded substantially to 29.2% and 7.5% respectively, MAT^3 maintained 57.5% success rate across all perturbations. This demonstrates that masked spatiotemporal representation learning enables robust adaptation to unseen conditions.

In summary, TaMeSo-bot with MAT^3 both with and without masking nearly completes the peg-in-hole tasks on seen pegs. Moreover, MAT^3 improves robustness to unseen pegs and conditions.

V. DISCUSSION

This section discusses the experimental results on the smoothness of retrieved actions and the learned embedding space, and how these relate to performance. We then discuss failure cases and outline future work.

TABLE I: Real-robot evaluation in **seen and unseen pegs**. Peg-in-hole task success rate (%) under four different starting positions.

Shape	Tactile Transformer	MAT^3 w/o Mask	MAT^3
<i>Seen</i>			
Square	8/40	36/40	35/40
Cyl. ($\varnothing 40$)	10/40	36/40	36/40
<i>Unseen</i>			
Cyl. ($\varnothing 30$)	12/40	33/40	35/40
Rectangle	6/40	26/40	32/40
Oval	7/40	28/40	34/40
Hexagon	5/40	30/40	36/40
Pentagon	5/40	25/40	33/40
Total (Seen)	18/80 (22.5%)	72/80 (90%)	71/80 (88.8%)
Total (Unseen)	35/200 (17.5%)	142/200 (71%)	170/200 (85%)

TABLE II: Real-robot evaluation with **unseen conditions** testing different starting positions, increased friction between the peg and the fixture surface, and a tilted grasping pose.

Condition	Tactile Transformer	MAT^3 w/o Mask	MAT^3
Unseen Starting Positions	2/40	9/40	20/40
Increased Friction	2/40	7/40	20/40
5° Tilted Grasp	5/40	19/40	29/40
Total	9/120 (7.5%)	35/120 (29.2%)	69/120 (57.5%)

A. Does MAT^3 help retrieve smoother actions?

Retrieved actions are not necessarily smooth; abrupt motions can cause task failures. Prior work [11] found action retrieval alone insufficient, potentially due to the lack of action smoothness, and therefore added a residual policy learned online via reinforcement learning. To assess smoothness with MAT^3 , Table III reports per-timestep differences between consecutive actions for MAT^3 and the Tactile Transformer. MAT^3 yields smaller differences for most peg shapes, indicating that it successfully retrieves smoother actions. The supplementary video also illustrates the abrupt motions produced by the Tactile Transformer.

B. Does MAT^3 's embedding space capture subtask-relevant information?

We hypothesize that our representation encodes subtask-relevant information and that retrieving actions sequentially for each subtask improves robustness. We qualitatively analyze the structure of MAT^3 's high-dimensional embedding space via visualization.

We build the retrieval database from 80% of the demonstrations and select representative query vectors from the remaining 20%. To visualize the high-dimensional structure while preserving neighborhood relations, we use a two-step procedure. First, we apply t-SNE to project offline state embeddings into a 2D space. Second, to relate high-dimensional proximity to subtasks, we overlay colormaps on

TABLE III: **Action smoothness comparison.** Each cell reports the **mean \pm std** of per-timestep action changes, computed across rollouts of the same peg shape. **Pos** is the Euclidean norm of position change and **Rot** is the angular change derived from the relative rotation. For each rollout, these increments are averaged over timesteps, then aggregated (mean \pm std) across all rollouts. **Lower values indicate smoother motion**, while greater values indicate jerky and discontinuous action. Bold marks the lower value between Transformer and MAT³ for each row and metric.

Shape	Tactile Transformer		MAT ³	
	Rot [°]↓	Pos [mm]↓	Rot [°]↓	Pos [mm]↓
<i>Seen</i>				
Square	0.3 \pm 0.2	2.2 \pm 1.1	0.2 \pm 0.1	1.7 \pm 0.6
Cyl. (\varnothing 40)	0.3 \pm 0.1	1.8 \pm 0.9	0.2 \pm 0.2	1.8 \pm 1.3
<i>Unseen</i>				
Cyl. (\varnothing 30)	0.2 \pm 0.2	1.8 \pm 1.6	0.2 \pm 0.2	1.7 \pm 1.0
Rectangle	0.4 \pm 0.2	2.5 \pm 0.9	0.2 \pm 0.1	1.7 \pm 0.6
Oval	0.3 \pm 0.2	2.4 \pm 0.8	0.2 \pm 0.1	1.6 \pm 0.5
Hexagon	0.2 \pm 0.1	2.0 \pm 0.7	0.1 \pm 0.1	1.4 \pm 0.8
Pentagon	0.2 \pm 0.1	2.0 \pm 0.5	0.1 \pm 0.1	1.8 \pm 0.6
Total (Seen)	0.3 \pm 0.2	2.0 \pm 1.0	0.2 \pm 0.1	1.7 \pm 1.0
Total (Unseen)	0.3 \pm 0.1	2.1 \pm 0.9	0.2 \pm 0.1	1.6 \pm 0.7

the t-SNE projection.

Fig. 6 shows colormaps of distances computed in the original embedding space between each projected state and query embeddings chosen for the subtasks *fit*, *align*, and *insert*. Query embeddings are taken from manually segmented sequences, using states at the end of each subtask. The t-SNE plots exhibit connected yet distinct regions for *fit*, *align*, and *insert*, indicating that MAT³ encodes subtask-related structure, which likely contributes to robustness.

C. Failure cases, limitations and future work

Failure cases include incorrect movement direction and partial insertion, where the peg either slips past the hole or gets caught, resulting in a gradual, viscous collapse of the end-effector pose.

As for limitations, this study focuses on peg-in-hole tasks as a common and active contact-rich manipulation task. Future work will consider various contact-rich manipulation tasks, such as grinding, scooping, and wiping. Skills in peg-in-hole, such as sliding, aligning, and inserting pegs while maintaining contact, can partially be shared with these tasks. Second, our approach assumes that the query and the embedding in the database are from reasonably close domains. While our method shows the adaptation to unseen pegs, for queries that fall completely outside the domain covered by the database, it will not extrapolate expected outputs. An interesting direction for future search would be to explore a method for transferring embedding spaces across domains, enabling more flexible adaptation to novel manipulation scenarios without requiring extensive data collection. In addition, this study uses an external motion tracker to capture the pose of the soft wrists. Future work could explore

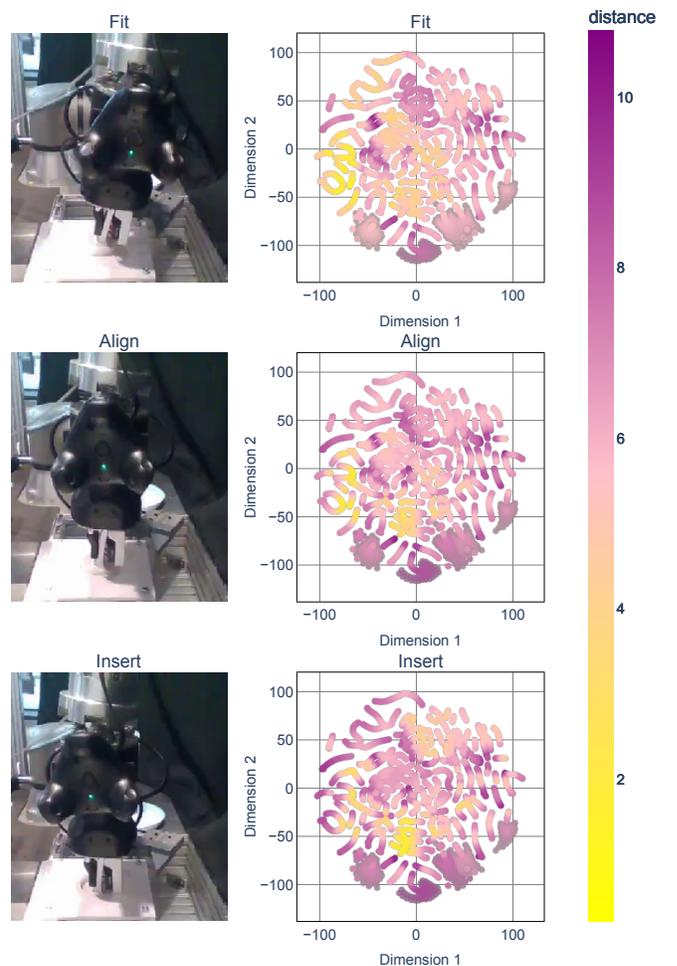


Fig. 6: **t-SNE visualization of the learned embedding space.** Colormaps indicate distance in the original high-dimensional space to representative query points for each subtask. **Right:** Embedding space t-SNE visualizations of MAT³, where points are colored by L2 distance to the query vector for each corresponding subtask. **Left:** Video frames corresponding to the *fit*, *align*, and *insert* subtasks.

incorporating visual inputs or modeling the entire system as a partially observable Markov decision process (POMDP) to realize a tracker-less system. Such an approach would simplify the system setup and enable broader reproduction of our method across diverse scenarios.

VI. CONCLUSION

We presented **TaMeSo-bot**, a system that integrates tactile memory with soft robotics for robust insertion tasks. Our Masked Tactile Trajectory Transformer (MAT³) learns spatiotemporal representations from multimodal sensor data while explicitly modeling spatial relationships across distributed taxel sensors. Real-robot experiments on peg-in-hole tasks demonstrate that MAT³ achieves robust generalization to unseen objects and conditions, significantly outperforming baselines without masked modeling or spatiotemporal structure. The integration of masked tactile representation

learning with physical compliance enables adaptation to novel scenarios. We believe the synergy of tactile memory and physical compliance opens promising new research directions, potentially enabling robots to achieve human-like dexterity and adaptability in complex manipulation tasks.

REFERENCES

- [1] A. Gallace and C. Spence, "The cognitive and neural correlates of tactile memory," *Psychological Bulletin*, vol. 135, no. 3, p. 380, 2009.
- [2] M. Suomalainen, Y. Karayiannidis, and V. Kyrki, "A survey of robot manipulation in contact," *Robotics and Autonomous Systems*, vol. 156, p. 104224, 2022.
- [3] S. Jin, X. Zhu, C. Wang, and M. Tomizuka, "Contact pose identification for peg-in-hole assembly under uncertainties," in *Proc. American Control Conference*, 2021, pp. 48–53.
- [4] F. von Drigalski, K. Tanaka, M. Hamaya, R. Lee, C. Nakashima, Y. Shibata, and Y. Ijiri, "A compact, cable-driven, activatable soft wrist with six degrees of freedom for assembly tasks," in *Proc. IROS*, 2020, pp. 8752–8757.
- [5] L. Johannsmeier, M. Gerchow, and S. Haddadin, "A framework for robot manipulation: Skill formalism, meta learning and adaptive control," in *Proc. ICRA*, 2019, pp. 5844–5850.
- [6] A. S. Morgan, Q. Bateux, M. Hao, and A. M. Dollar, "Towards generalized robot assembly through compliance-enabled contact formations," in *Proc. ICRA*, 2023, pp. 8010–8016.
- [7] R. S. Johansson and J. R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks," *Nature Reviews Neuroscience*, vol. 10, no. 5, pp. 345–359, 2009.
- [8] "Contactile Accessed on: Apr. 21, 2025," <https://contactile.com/>.
- [9] N. Di Palo and E. Johns, "On the effectiveness of retrieval, alignment, and replay in manipulation," *Proc. ICML*, vol. 9, no. 3, pp. 2032–2039, 2024.
- [10] I. Guzey, B. Evans, S. Chintala, and L. Pinto, "Dexterity from touch: Self-supervised pre-training of tactile representations with robotic play," in *Proc. CoRL*, 2023, pp. 3142–3166.
- [11] K. Yu, Y. Han, Q. Wang, V. Saxena, D. Xu, and Y. Zhao, "Mimictouch: Leveraging multi-modal human tactile demonstrations for contact-rich manipulation," in *Proc. CoRL*, 2024.
- [12] Y. Chen, M. Van der Merwe, A. Sipos, and N. Fazeli, "Visuo-tactile transformers for manipulation," in *Proc. CoRL*, 2023, pp. 2026–2040.
- [13] P. Wu, A. Majumdar, K. Stone, Y. Lin, I. Mordatch, P. Abbeel, and A. Rajeswaran, "Masked trajectory models for prediction, representation, and control," in *Proc. ICML*, 2023, pp. 37 607–37 623.
- [14] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL*. Minneapolis, Minnesota: Association for Computational Linguistics, 2019, pp. 4171–4186.
- [15] J. Qu, B. Mao, Z. Li, Y. Xu, K. Zhou, X. Cao, Q. Fan, M. Xu, B. Liang, H. Liu *et al.*, "Recent progress in advanced tactile sensing technologies for soft grippers," *Advanced Functional Materials*, vol. 33, no. 41, p. 2306249, 2023.
- [16] N. F. Lepora, C. Ford, A. Stinchcombe, A. Brown, J. Lloyd, M. G. Catalano, M. Bianchi, and B. Ward-Cherrier, "Towards integrated tactile sensorimotor control in anthropomorphic soft robotic hands," in *Proc. ICRA*, 2021, pp. 1622–1628.
- [17] O. Shorthose, A. Albin, L. He, and P. Maiolino, "Design of a 3d-printed soft robotic hand with integrated distributed tactile sensing," *Proc. ICML*, vol. 7, no. 2, pp. 3945–3952, 2022.
- [18] Z. Zhou, R. Zuo, B. Ying, J. Zhu, Y. Wang, X. Wang, and X. Liu, "A sensory soft robotic gripper capable of learning-based object recognition and force-controlled grasping," *IEEE Transactions on Automation Science and Engineering*, vol. 21, no. 1, pp. 844–854, 2022.
- [19] C. Lu, K. Tang, M. Yang, T. Yue, H. Li, and N. F. Lepora, "Dexitac: Soft dexterous tactile gripping," *IEEE/ASME Transactions on Mechatronics*, vol. 30, no. 1, pp. 333–344, 2024.
- [20] N. Zhang, J. Ren, Y. Dong, X. Yang, R. Bian, J. Li, G. Gu, and X. Zhu, "Soft robotic hand with tactile palm-finger coordination," *Nature Communications*, vol. 16, no. 1, p. 2395, 2025.
- [21] M. A. Lin, E. Reyes, J. Bohg, and M. R. Cutkosky, "Whisker-inspired tactile sensing for contact localization on robot manipulators," in *Proc. IROS*, 2022, pp. 7817–7824.
- [22] M. S. Sofla, S. Vayakkattil, and M. Calisti, "Haptic localization with a soft whisker from moment readings at the base," *Soft Robotics*, vol. 11, no. 4, pp. 539–549, 2024.
- [23] H. Li, C. Xing, S. Khan, M. Zhong, and M. R. Cutkosky, "Whisker-inspired tactile sensing: A sim2real approach for precise underwater contact tracking," *Proc. ICML*, 2025.
- [24] U. Yoo, N. Dennler, E. Xing, M. Matorić, S. Nikolaidis, J. Ichnowski, and J. Oh, "Soft and compliant contact-rich hair manipulation and care," in *Proc. HRI*, 2025, pp. 610–619.
- [25] C. Zhao, C. Jiang, L. Luo, S. Yuan, Q. Chen, and H. Yu, "Learning thin deformable object manipulation with a multi-sensory integrated soft hand," *IEEE Transactions on Robotics*, vol. 41, pp. 4606–4621, 2025.
- [26] J. Royo-Miquel, M. Hamaya, C. C. Beltran-Hernandez, and K. Tanaka, "Learning robotic assembly by leveraging physical softness and tactile sensing," in *Proc. IROS*, 2023, pp. 6469–6476.
- [27] O. Azulay, D. M. Ramesh, N. Curtis, and A. Sintov, "Visuotactile-based learning for insertion with compliant hands," *Proc. ICML*, vol. 10, no. 4, pp. 4053–4060, 2025.
- [28] Q. Li, O. Kroemer, Z. Su, F. F. Veiga, M. Kaboli, and H. J. Ritter, "A review of tactile information: Perception and action through touch," *IEEE Transactions on Robotics*, vol. 36, no. 6, pp. 1619–1634, 2020.
- [29] C. Higuera, A. Sharma, C. K. Bodduluri, T. Fan, P. Lancaster, M. Kalakrishnan, M. Kaess, B. Boots, M. Lambeta, T. Wu *et al.*, "Spash: Self-supervised touch representations for vision-based tactile sensing," in *Proc. CoRL*, 2024.
- [30] J. Zhao, Y. Ma, L. Wang, and E. Adelson, "Transferable tactile transformers for representation learning across diverse sensors and tasks," in *Proc. CoRL*, 2024.
- [31] A. Garcia-Garcia, B. S. Zapata-Impata, S. Orts-Escolano, P. Gil, and J. Garcia-Rodriguez, "Tactilegnn: A graph convolutional network for predicting grasp stability with tactile sensors," in *International Joint Conference on Neural Networks*, 2019, pp. 1–8.
- [32] L. Yang, B. Huang, Q. Li, Y.-Y. Tsai, W. W. Lee, C. Song, and J. Pan, "Tacgnn: Learning tactile-based in-hand manipulation with a blind robot using hierarchical graph neural network," *Proc. ICML*, vol. 8, no. 6, pp. 3605–3612, 2023.
- [33] S. Kulkarni, S. Funabashi, A. Schmitz, T. Ogata, and S. Sugano, "Tactile object property recognition using geometrical graph edge features and multi-thread graph convolutional network," *Proc. ICML*, 2024.
- [34] S. Dong, D. K. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez, "Tactile-rl for insertion: Generalization to objects of unknown geometry," in *Proc. ICRA*, 2021, pp. 6437–6443.
- [35] T. Ueno, S. Funabashi, H. Ito, A. Schmitz, S. Kulkarni, T. Ogata, and S. Sugano, "Multi-fingered dragging of unknown objects and orientations using distributed tactile information through vision-transformer and lstm," in *Proc. IROS*, 2024, pp. 7445–7452.
- [36] M. Du, S. Nair, D. Sadigh, and C. Finn, "Behavior retrieval: Few-shot imitation learning by querying unlabeled datasets," *RSS*, 2023.
- [37] S. Nasiriany, T. Gao, A. Mandlekar, and Y. Zhu, "Learning and retrieval from prior data for skill-based imitation learning," in *Proc. CoRL*, 2023, pp. 2181–2204.
- [38] S. Kuroki, M. Nishiura, and T. Kozuno, "Multi-agent behavior retrieval: Retrieval-augmented policy training for cooperative push manipulation by mobile robots," in *Proc. IROS*, 2024.
- [39] P. C. Humphreys, A. Guez, O. Tieleman, L. Sifre, T. Weber, and T. P. Lillicrap, "Large-scale retrieval for reinforcement learning," in *Proc. NeurIPS*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.
- [40] G. Papagiannis, N. D. Palo, P. V. itiello, and E. Johns, "R+x: Retrieval and execution from everyday human videos," in *Proc. ICRA*, 2025.
- [41] N. Di Palo and E. Johns, "On the effectiveness of retrieval, alignment, and replay in manipulation," *Proc. ICML*, 2024.
- [42] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proc. CVPR*, 2017, pp. 652–660.
- [43] J. Lee, Y. Lee, J. Kim, A. Kosiorek, S. Choi, and Y. W. Teh, "Set transformer: A framework for attention-based permutation-invariant neural networks," in *Proc. ICML*, 2019, pp. 3744–3753.
- [44] Y. A. Malkov and D. A. Yashunin, "Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 4, pp. 824–836, 2018.