

Multi-Resolution Tactile Imitation Learning for Contact-Rich Robotic Manipulation

Rickmer Krohn^{*,1,3,4}, Erik Helmut², Niklas Funk²,
Jan Peters^{2,3,4}, Vignesh Prasad^{1,3,4} and Georgia Chalvatzaki^{1,3,4}

¹ Interactive Robot Perception & Learning, TU Darmstadt

² Intelligent Autonomous Systems, TU Darmstadt

³ Hessian AI, ⁴ Robotics Institute Germany

*rickmer.krohn@tu-darmstadt.de

Abstract: Touch sensing is beneficial for solving a wide variety of manipulation tasks. While there exists a wide range of tactile sensors with different properties, exploiting the fusion of multiple heterogeneous tactile sensors to improve manipulation learning remains underexplored. We present Multi-Resolution Tactile Sensing (MiTaS), a representation framework that leverages multiple tactile sensors operating at different temporal resolutions in order to solve complex contact-rich manipulation tasks. We propose a novel architecture using modality-specific convolutional stems and transformer-based fusion that effectively fuses information from an RGB camera stream, a vision-based GelSight Mini sensor and a high-frequency event-based Evetac sensor. This multi-sensor representation then conditions a flow-matching policy for solving downstream tasks. Experimental results across five contact-rich manipulation tasks demonstrate the effectiveness of multi-resolution tactile features in imitation learning. MiTaS achieves an average success rate of 80 %, while vision-only (31 %) and visual-tactile (54 %) baselines cannot solve the task reliably. Co-training a visuo-tactile model with multi-tactile data boosts performance by over 10 % in certain tasks, without having access to the Evetac sensor during policy evaluation. A detailed sensor-reading and attention analysis reveals the importance of different sensors throughout task execution, validating our multi-resolution tactile sensing approach.

Project Page: <http://mitas-touch.github.io>.




Figure 1: MiTaS combines an RGB camera (blue), a GelSight Mini (red) and an event-based Evetac sensor (yellow) mounted on the gripper. Our method fuses these multi-resolution tactile streams to condition a flow-matching policy, enabling improved control in contact-rich manipulation tasks such as key insertion.

1 Introduction

Humans rely on multi-modal touch sensing to accomplish delicate manipulation tasks, seamlessly integrating information from diverse mechanoreceptors operating at different temporal and spatial resolutions [1]. Similarly, robotic systems can benefit from heterogeneous tactile sensing to understand objects, tasks, and contact dynamics, enabling precise manipulation capabilities.

RGB vision-based tactile sensors such as GelSight [2, 3, 4] have shown remarkable results in manipulation by providing high-resolution spatial information about local contact geometry, deformation, and surface structure [5, 6, 7, 8, 9, 10, 11]. However, such frame-based tactile sensing is less suited to capturing very fast contact changes, such as impact, incipient slip, high-frequency vibration. These changes may occur between frames or be smoothed out by conventional image-based processing. Recent advances in tactile sensing technology have also introduced event-based tactile cameras such as Evetac [12] or GelEvent [13], which offer high temporal resolution to capture rapid contact dynamics that may otherwise go unnoticed by traditional vision-based sensors. These high-frequency signals can detect critical events such as early contact, incipient slip, vibrations during insertion, and subtle variations during sliding contact.

Given the complementary strengths of classic vision-based and event-based sensors, a crucial question arises: can robotic systems benefit from leveraging heterogeneous tactile sensing capabilities to better understand intricate contact dynamics and improve manipulation learning?

In this work, we present **Multi-Resolution Dynamic Tactile Sensing (MiTaS)**, a comprehensive representation framework to condition a flow matching imitation learning policy to leverage multiple tactile sensors, operating at different temporal resolutions, to solve complex manipulation tasks. Our key insight is that different tactile modalities capture complementary aspects of contact-rich interactions: standard vision-based sensors provide high spatial resolution of contact geometry, while event-based sensors capture high-frequency temporal dynamics. We introduce a novel architecture that extracts and fuses heterogeneous tactile features through modality-specific CNN-stems and transformer-based fusion.

Our experimental results demonstrate that utilizing multiple tactile sensors enables the policy to achieve superior performance across five challenging contact-rich manipulation tasks. The MiTaS framework achieves an average success rate of 80%, whereas vision-only models are not able to solve all tasks, showcasing the effectiveness of incorporating multiple heterogeneous tactile sensors across dynamic and fine-grained tasks. An analysis of attention scores throughout task execution reveals how sensor importance shifts during different phases of manipulation, validating our multi-resolution sensing hypothesis. To the best of our knowledge, this is the first work to combine a classic vision-based tactile sensor with an event-based tactile sensor, thereby showing how complementary tactile features enable complex contact-rich manipulation.

Our contributions are the following: (1) A novel architecture for extracting and fusing heterogeneous tactile sensor readings operating at different temporal resolutions; (2) A multi-tactile co-training scheme that improves visual-tactile policy performance on contact-rich tasks where reactivity to small perturbations is key; (3) A multi-resolution tactile dataset combining egocentric RGB, vision-based, and event-based tactile sensing across five challenging contact-rich manipulation tasks.

2 Related Work

Tactile Sensing for Robotic Manipulation. Touch feedback has become increasingly important for contact-rich robotic manipulation [14], particularly in settings where local contact state cannot be reliably inferred from vision alone [5, 15, 16, 17, 18, 19, 7, 20]. Different families of tactile sensors have been successfully used to improve manipulation policies. Vision-based tactile sensors, such as GelSight [2], provide dense spatial information about contact geometry, while event-based tactile sensors such as Evetac [12] and GelEvent [13] capture high-frequency contact changes.

As tactile sensing has gained traction, reactive diffusion policies have incorporated visual-tactile feedback for contact-rich manipulation [21], demonstrating the value of multi-modal conditioning. Large-scale datasets such as Touch100k [22] have enabled touch-centric multimodal representation learning by providing diverse vision-language-touch data. Recent work has explored representations that go beyond pixel-level tactile information [23] and developed tactile-conditioned diffusion policies for force-aware manipulation [7]. Pretraining approaches such as ViTaL [24] have demonstrated that visuo-tactile pretraining benefits both tactile and non-tactile manipulation policies. The principle that "touch begins where vision ends" [25] has motivated generalizable policies for contact-rich manipulation. Self-supervised contrastive pretraining [26] has further advanced multimodal visual-tactile representation learning. Deep reinforcement learning with visuotactile feedback [8] has enabled learning of complex manipulation policies, while safe self-supervised learning [27] has made real-world industrial insertion feasible. Recent vision-tactile-language-action models [28] incorporate preference learning for insertion tasks, and world models such as DreamTacVLA [29] enable learning to predict future tactile feedback. Unified representations across static and dynamic tactile sensors [30] have further demonstrated the potential of multi-sensor tactile learning. Along similar lines, transferable tactile transformers [31] have shown promise for learning representations across diverse sensors and tasks. Despite these significant advancements in tactile-driven manipulation, most current frameworks rely on a single touch sensor modality. Drawing inspiration from the multimodal nature of human touch, this work bridges this gap by exploring how fusing two distinct tactile sensors with complementary temporal resolutions can improve the robustness of learning robotic manipulation.

Multi-Sensory Learning for Manipulation. The integration of vision and touch has been explored extensively for improving manipulation capabilities [16, 19, 32, 33]. Multi-sensor fusion with stage guidance has been applied to complex tasks such as pouring and peg insertion [34]. The combination of vision and tactile via CNN-stems and masked autoencoding improves visuotactile and vision policies [35]. Sparsh-X [23] is training a representation on the multimodal Digit360 [36] sensor via Self-Supervised Learning. The Fusion of Vision, Audio and Touch via a self-attention layer improves single-modality performance [18, 37]. Visual-force-tactile fusion [38] has enabled gentle intricate insertion tasks by leveraging complementary information from multiple sensing modalities. Multi-modal human tactile demonstrations [39] have been leveraged to improve contact-rich manipulation through imitation learning. Nevertheless, these approaches focus on cross-modal integration (e.g., vision plus audio) rather than on interplay among tactile-sensing modalities. Despite the promise of multi-sensory learning, fusing co-located tactile modalities with complementary temporal resolutions to condition manipulation policies remains largely unexplored, a gap that MiTaS directly addresses.

3 Multi-Resolution Tactile Sensing

In this section, we present Multi-Resolution Tactile Sensing (MiTaS), a novel framework that leverages multiple heterogeneous tactile sensors operating at different temporal resolutions to learn manipulation policies through imitation. Our proposed approach consists of three main components: (1) a novel architecture utilizing modality-specific CNN stems that encode diverse sensor inputs into unified token representations, (2) a flow-matching policy for action generation and visuomotor control conditioned on these tokens, and (3) a multitactile co-training scheme that leverages the all available sensors during training to improve downstream policy performance when only a subset of sensors is available at inference time. To the best of our knowledge this is the first work to combine a classic vision-based tactile sensor with an event-based tactile sensor, showing how complementary tactile features enable complex contact-rich manipulation. Figure 2 provides an overview of the MiTaS encoder architecture and policy training.

For training the multi-sensory visuomotor policies we consider three complementary sensing modalities: egocentric RGB from a wrist-mounted camera, GelSight Mini tactile images [2], and high-rate Evetac tactile frames [12]. At each timestep, the observation is $o_t = [O_t^{\text{vision}}, O_t^{\text{gelsight}}, O_t^{\text{evetac}}]$, where $O_t^{\text{vision}} \in \mathbb{R}^{2 \times 128 \times 128}$ is recorded at 25 Hz and captures the scene and hand-object con-

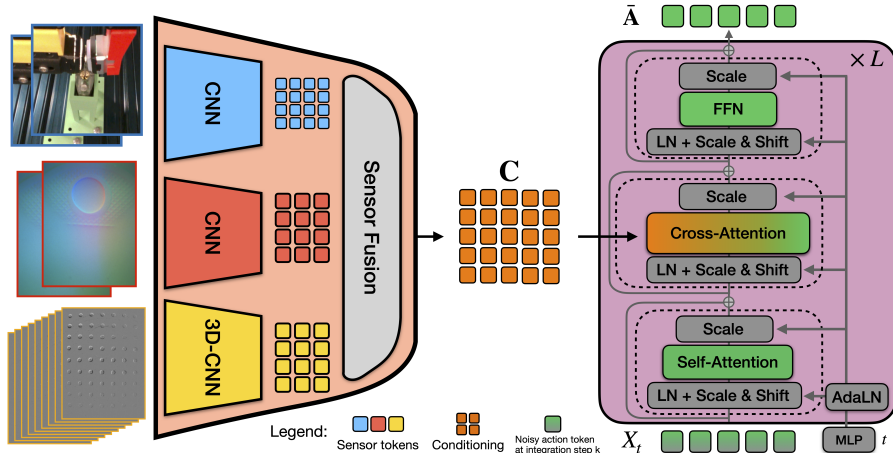


Figure 2: Overview of the MiTaS architecture: Modality-specific CNN stems encode Vision, GelSight, and Evetac sensors into token embeddings. These tokens are fused through transformer-based attention mechanisms and form the condition for a flow matching policy. The policy does not receive the robot state and must infer the delta position prediction solely from sensor readings.

figuration, $O_t^{\text{gelsight}} \in \mathbb{R}^{2 \times 120 \times 160}$ is recorded at 25 Hz and captures local contact geometry, and $O_t^{\text{evetac}} \in \mathbb{R}^{16 \times 120 \times 160}$ stacks 16 event-based frames per timestep from the 200 Hz sensor to capture fast contact dynamics. Vision and GelSight observations are normalized to $[0, 1]$, while Evetac is normalized to $[-\frac{1}{2}, \frac{1}{2}]$ so the default grey frame corresponds to 0 (no event activity) rather than 0.5. These input observations are then passed into our proposed MiTaS Architecture, and we subsequently leverage a flow-matching generative model [40, 41] for predicting a chunk of end-effector delta actions. We want to highlight that since the policy is neither conditioned on robot state nor expressed in absolute position space, it must infer the appropriate relative motion directly from high-dimensional multimodal sensory input. In the following, we first detail the MiTaS encoding architecture, then the policy parameterization for action generation, and finish with our proposed co-training procedure.

3.1 Visual Multi-Tactile Encoding Architecture

Extracting and fusing features from heterogeneous sensor modalities is essential to effectively learn multimodal imitation learning policies. We propose a transformer-based architecture to handle multimodal data during encoding and using cross-attention to condition the policy. Cross-attention allows for efficient integration of high-dimensional data, without introducing a bottleneck on the perception side. The Encoder consists of sensor stems, positional / modality encoding, and sensor fusion. An overview of the Architecture is shown in Figure 2.

Sensor Stems: Each modality is tokenised by a dedicated CNN-stem to extract a meaningful sensor-specific representation. The wrist-camera stacks two RGB frames along the channel axis and applies four strided 2D convolutions, producing a 16×16 grid of $D=256$ -dimensional tokens. The GelSight stem mirrors this design at the rectangular 120×160 resolution and yields a 12×16 token grid. Evetac’s 16-frame stack is processed by a four-layer 3D-convolutional stem that jointly compresses time and space into a 12×16 grid, reflecting the higher temporal density of that modality. The 3D-CNN uses temporal kernels of 3–3–3–2 and strides of 2 in time and space through the first three stages, then a final $(2 \times 4 \times 5)$ convolution layer that collapses time to form the token features. The embeddings of the tactile modalities follow the spatial dimension of the sensor observation and allow associations between the different tactile readings later on. Furthermore, early convolutions result in overlapping receptive fields, providing spatially coherent features and have been shown to stabilize training [42].

Positional and modality encoding: Each sensor stem outputs a grid of D -dimensional tokens per sensor; after flattening, every token z_i is augmented with a learned position vector $p_{g(i)}$ that encodes its location within that sensor’s token grid, and a modality-specific embedding $e_{s(i)}$ where

$s(i) \in \{\text{vision, gelsight, evetac}\}$ identifies the source sensor, i.e., $z'_i = z_i + \mathbf{p}_{g(i)} + e_{s(i)}$. This separates *where* a token lies in the spatial layout from *which* modality produced it. Per-token positional encodings and modality-specific segment embeddings are learned jointly with the policy.

Sensor Fusion: The augmented visual multi-tactile token sequence is then processed through a transformer encoder [43] with multi-head self-attention, MLP-layers, pre-normalization and residual connections. The fusion transformer allows full self- and cross-attention of sensors, leading to rich multisensory features forming the conditioning for the policy.

3.2 Action Generation through Flow Matching

Next, we utilize the visual multi-tactile embeddings for robot action generation. Following recent advances in learning robotic manipulation policies from expert demonstrations [44, 45], we leverage a conditional flow matching generative model for policy learning [40]. We define our policy output as a fixed-length sequence of future robot commands $\mathbf{A}_t = [\mathbf{a}_t, \mathbf{a}_{t+1}, \dots, \mathbf{a}_{t+H-1}] \in \mathbb{R}^{H \times d_a}$, with horizon H and action dimension d_a . The individual actions are *delta* commands in the robot local frame, i.e., $\mathbf{a} = \Delta \mathbf{p} = [\Delta x, \Delta y, \Delta z, \Delta \psi] \in \mathbb{R}^4$. This formulation controls translation and the yaw rotation around the z-axis, while keeping pitch and roll fixed. The control loop executes only the first $n_{\text{act}} \leq H$ predicted steps before replanning, improving reactivity while retaining a temporally coherent training target. The policy head is a Diffusion-Transformer (DiT) that regresses a time-conditioned velocity field along a straight noise-to-action path.

Flow-matching training and inference: During training, we start with sampling an initial random noise vector $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and we also sample a desired data endpoint from our available demonstrations, i.e., $\mathbf{x}_1 = \bar{\mathbf{A}}$. Following the conditional flow-matching paradigm [41], we define a straight-line probability path \mathbf{x}_t and its corresponding constant target velocity field \mathbf{u}_t for $t \in [0, 1]$ as: $\mathbf{x}_t = t \cdot \mathbf{x}_1 + (1-t) \cdot \mathbf{x}_0$, and $\mathbf{u}_t = d\mathbf{x}_t/dt$. The training objective is to optimize the network parameters θ such that its output $\hat{\mathbf{v}}_\theta(\mathbf{x}_t, t, \mathbf{C})$ that is conditioned on the current intermediate action state \mathbf{x}_t , the flow time t , and the vision and multitactile tokens \mathbf{C} accurately regresses this target velocity field, i.e., we minimize mean-squared error $\|\hat{\mathbf{v}}_\theta(\mathbf{x}_t, t, \mathbf{C}) - \mathbf{u}_t\|^2$ over all the flow time steps across all demonstrations. During inference, we generate actions by integrating the learned velocity field via explicit Euler integration starting from an initial noise vector, $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and iterate $\mathbf{x}_{t+\Delta t} \leftarrow \mathbf{x}_t + \Delta t \hat{\mathbf{v}}_\theta(\mathbf{x}_t, t, \mathbf{C})$ for $t \in [0, 1)$ with a fixed $\Delta t = 1/K$.

DiT action head: The denoiser network ultimately outputs $\hat{\mathbf{v}}_\theta(\mathbf{x}_t, t, \mathbf{C})$, thereby operating on an intermediate action state $\mathbf{x}_t \in \mathbb{R}^{H \times d_a}$. Each horizon slot is linearly projected to a D -dimensional token and summed with a learned positional embedding over the chunk, yielding $\mathbf{X}_t \in \mathbb{R}^{H \times D}$. A stack of L DiT blocks updates \mathbf{X}_t while conditioning on multisensory tokens $\mathbf{C} \in \mathbb{R}^{N_{\text{cond}} \times D}$ produced by the perception encoder; queries come from the action tokens and keys/values from \mathbf{C} , and we use cross-attention without an intermediate projection since both streams share dimension D . Each block applies (i) multi-head self-attention over the H action tokens, (ii) multi-head cross-attention to \mathbf{C} , and (iii) a token-wise MLP, using residual connection throughout. We condition each block on t via AdaLN-Zero [46, 47] to provide a strong, stable time-dependent modulation while initializing the network near an identity mapping (details in Appendix A.3). The last layer maps the updated tokens back to the final output $\hat{\mathbf{v}}_\theta(\mathbf{x}_t, t, \mathbf{C}) \in \mathbb{R}^{H \times d_a}$.

3.3 Multi-Tactile Co-Training

Recent work have shown that training a policy on multisensory data can improve vision-based policies, either with pretraining [35, 24, 48] or masking [20]. Building on these insights, we introduce an asymmetric co-training scheme designed to train a visual-tactile policy (utilizing vision and Gel-Sight) while exploiting the high-frequency Evetac sensor exclusively during the training phase. Each training batch is constructed by sampling half of the sequences with Evetac observations and the remaining half without. This stochastic co-training procedure allows the high-frequency temporal dynamics captured by Evetac to shape the representation space of the sensor-fusion network. Crucially, because the model must simultaneously learn to generate accurate actions in the absence of Evetac, this procedure effectively regularizes the standard visual-tactile features to latently encode

and compensate for the missing features extracted from Evetac observations. The Evetac encoder is omitted entirely during deployment, leaving a streamlined (co-trained) vision-and-gelsight policy that benefits from cross-modal feature alignment established during the training phase.

4 Experimental Evaluation

We evaluate MiTaS on five challenging contact-rich manipulation tasks that require tactile feedback: Gear assembly, Board Wiping, Lamp Installation, Key in Lock, and Lightbulb Connection. Our experiments investigate: (1) overall performance compared to a strong multi-sensory baseline and vision-only approaches, (2) the contribution of individual sensors through ablation studies, (3) impact of multi-tactile co-training, and (4) the dynamic importance of different sensors throughout task execution.

Hardware Platform and Data Collection: This Project uses a 7-DoF Franka Panda robot with an attached Robotis RH-P12-RN gripper. All sensors are mounted on the gripper with 3D-printed holders. Vision is obtained by a wrist-mounted Intel RealSense D435 camera. Gelsight Mini and Evetac are attached to one finger each and form the parallel gripper. The robot is moved with a Cartesian impedance controller. We collect 30 expert demonstrations by teleoperating the robot via a Spacemouse. The control frequency is 15Hz, and sensor frames are uniformly sampled in a 0.06 s window to form the sensor history per observation step.

Baseline and Sensor ablation: We compare the MiTaS architecture against Sparsh-X [23], a multimodal tactile transformer baseline. Sparsh-X uses 8 unimodal and 4 fusion transformer layers, in which attention bottlenecks [49] are used as a fusion mechanism instead of self-attention. All models are trained end-to-end with the policy objective using the same DiT action-head. Therefore, the models only differ in architecture and sensor settings to shape the multisensory conditioning of the flow matching policy. For the sensor ablation, we simply omit the encoding of the respective sensor(s). Vision-only models resemble the Vision Transformer architecture [50], following the sensor stems from MiTaS and Sparsh-X, we call them ViT-CNN and ViT.

4.1 Task Setups

Gear Assembly [51] requires inserting a cylindrical gear between two placed gears. The gear must be aligned in position and orientation to seat on the holder. The gear-teeth can cause it to jam, which is resolved by adjusting orientation and applying fine-grained position changes. **Board Wiping** entails wiping a marked line from a fixed board. Contact must be maintained along the line throughout the wipe. Insufficient pressure fails to remove the mark, while excessive pressure tilts or wedges the sponge and breaks the planar contact. **Lamp Installation** [52] requires to align a 3D-printed bulb in a socket in order to screw its threaded base into it. Once the threads are engaged, the bulb is rotated about its vertical axis to drive the threads. Loss of axial alignment can cause it to jam before seating. **Key in Lock** involves inserting a key fully into a real lock. The key must align with the keyway in both position and orientation. The long and tight insertion path causes even small angular deviations to bind the key against the internal walls. Precise alignment must hold across the full movement until the key is seated. **Lightbulb Connection** consists of grasping a GU10 LED bulb and connecting it to its socket. The two posts at the base of the bulb must align with the socket slots to enter, after which the bulb is rotated about its vertical axis until the post heads slide into their locking slots. See Figure 3 for an overview of the tasks and Figure 6 (Appendix) for full task rollouts.




Figure 3: Representative images of the five contact-rich manipulation tasks in order from left to right: Gear Assembly, Board Wiping, Lamp Installation, Key in Lock, and Lightbulb Connection.

4.2 Results

We first evaluate overall task performance against the multisensory baseline and vision-only models. Followed by a tactile sensor ablation study to showcase the benefits of multi-resolution tactile sensing compared to visual-tactile models. To isolate the effect of sensor encoding, all models share the same flow-matching policy head and are trained on the same demonstration data, differing only in which and how the sensor inputs are encoded into conditioning tokens. Lastly, we evaluate our multi-tactile co-training scheme against the sensor ablation from above. We evaluate every policy 20 times and report the success rate. A table with all results can be found in the Appendix A.8.




Figure 4: MiTaS outperforms the multimodal baseline Sparsh-X and two vision-only ViT-baselines in every task. Vision-only models failed, due to occlusion and complex contact dynamics, highlighting the need for tactile sensing in contact-rich manipulation.

Figure 4 shows the success rates of MiTaS, Sparsh-X, and both vision-only variants ViT and ViT-CNN over five manipulation tasks. MiTaS, with an average success rate of 80%, is able to solve all tasks consistently, while Sparsh-X (avg. 54%), with the same sensor setting, is not reliable in every task. Remarkably, Sparsh-X inserted the key multiple times but never completed the task, getting stuck in the keyway. Both vision-only models (ViT 31%; ViT-CNN 26%) have a low performance across all tasks and are not able to solve Key in Lock and Lightbulb, due to tight insertions dynamics and visual occlusions. Multi-resolution tactile feedback paired with suitable encoding enables a high success rate with only 30 demonstrations in these complex contact-rich manipulation tasks.

Sensor Ablation Study. We compare the full sensor setting against the vision + gelsight setting. Additionally, we have trained the MiTaS model on Vision and Evetac investigating its performance without Gelsight grounding. Table 1 shows the sensor ablation results. Visual multi-tactile models achieve the highest success rate regardless of model design, confirming that heterogeneous tactile sensors provide complementary information required to solve contact-rich manipulation task.

Still, simple vision-tactile policy are able to solve most of the task over half of the time. Evetac as the sole tactile sensor provides useful information in certain task, but fails completely when object in hand location is needed. Gelsight, on the other hand, can align the gear and key from various initial states (see Appendix A.7).

Multi-tactile Co-Training. Table 2 shows the impact of our proposed multi-tactile co-training, leveraging the Evetac sensor during training and not requiring it during policy evaluation. It boosted MiTaS performance in 3/5 tasks, while being helpful in 2/5 tasks for Sparsh-X. In tasks, like Key in Lock, the co-training scheme lowers performance as it disturbs Gelsight’s features for in-hand localization. In Board

Table 1: Policy success rate (%).

Task	MiTaS (Ours)			Sparsh-X	
	V+G+E	V+G	V+E	V+G+E	V+G
Gear	90%	45%	0%	50%	25%
Board	90%	70%	80%	85%	50%
Lamp	80%	65%	65%	70%	55%
Key	75%	55%	0%	10%	20%
Lightbulb	65%	20%	40%	55%	15%
Avg.	80%	51%	37%	54%	33%

Table 2: Co-training ablation at V+G (success in %).

Task	MiTaS (Ours)			Sparsh-X		
	V+G	+Co-train	Δ	V+G	+Co-train	Δ
Gear	45%	55%	+10	25%	15%	-10
Board	70%	85%	+15	50%	85%	+35
Lamp	65%	60%	-5	55%	70%	+15
Key	55%	20%	-35	20%	5%	-15
Lightbulb	20%	45%	+25	15%	10%	-5
Avg.	51%	53%	+2	33%	37%	+4

Wiping, both visual-tactile models benefit from the Evetac co-training, aligning with the insights from the sensor ablation (MiTaS V+E) that Evetac features are beneficial in this dynamic tasks.

4.3 Attention Analysis with Sensor Visualization

To understand how MiTaS uses the different sensors over the course of a task, we visualize the cross-attention from the executed action token to each sensor modality during a successful episode of the Lamp Installation task.

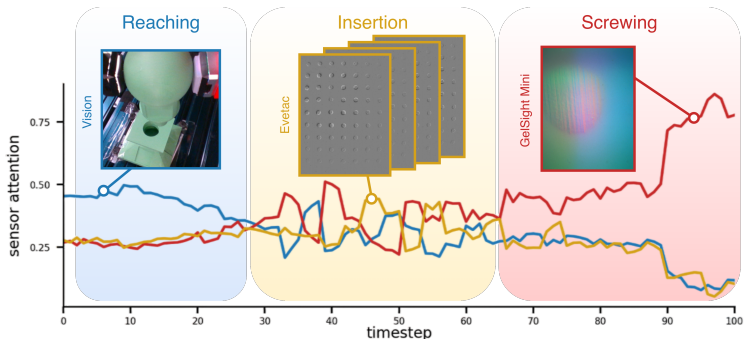


Figure 5 shows how the policy attends on vision

Figure 5: Attention analysis with sensor readings in Lamp Installation.

during reaching, on Gelsight for the prolonged screwing motion and on all sensors during insertion. The improved performance and attention show that evetac is not driving the representation, but offers relevant features during crucial stages of the task. We observed recovery behaviour, when it was jammed during Gear assembly and resolving ambiguity in the Lighbulb Installation task, where Gelsight could not differentiate between hovering closely over the socket or being fully inserted.

5 Limitations

While MiTaS demonstrates strong performance across our contact-rich task suite, several limitations remain. First, our investigation of multi-resolution fusion focuses exclusively on a specific pair of optical tactile sensors: the frame-based GelSight and the event-based Evetac. While this setup allowed for an exact and detailed analysis of complementary spatial and temporal resolutions, further research is required to evaluate how our architectural framework scales when integrating other heterogeneous sensor pairs or alternative tactile technologies. Second, our data collection pipeline relies on teleoperation via a Spacemouse, which fundamentally constrains scalability. Transitioning to alternative demonstration sources, such as human video or markerless hand tracking without a physical robot in the loop, could drastically scale up data collection. Furthermore, our policy is currently limited to 4-DoF actions by choice, though the framework itself imposes no fundamental limit on action dimensionality. While a 4-DoF space sufficed for our target contact-rich tasks, scaling to full 6-DoF control would likely require incorporating proprioceptive robot state information, which is currently omitted from our policy’s observation space. Finally, the 15 Hz replanning frequency of our control loop places an upper bound on the framework’s reactivity. Integrating our multi-resolution sensing framework with recent advances in highly reactive policy architectures offers a promising direction to further improve policy performance.

6 Conclusion

In this work, we propose Multi-Resolution Tactile Sensing (MiTaS), a novel imitation learning framework that leverages multiple tactile sensors operating at different temporal resolutions to solve complex manipulation tasks. Specifically, MiTaS encodes vision, GelSight, and Evetac frames via CNN-stems and transformer fusion to condition a flow matching policy. Vision combined with heterogeneous tactile features obtained at low and high frequencies enables fine-grained manipulation, even under heavy occlusion, in complex contact-rich manipulation tasks. MiTaS achieves 80%, while the best baselines have a success rate of 54%. Vision-only policies fail completely in two tasks, highlighting the need for tactile sensors in contact-rich tasks. Furthermore, we have shown how multisensory co-training can boost a classic visual-tactile policy in certain task. In future work, we want to investigate event-based tactile features for representation learning on a larger scale and a co-training regime to improve widespread vision-based tactile sensors for manipulation, next to more reactive policies using high-frequency tactile sensors.

Acknowledgments

This research is funded by the German Research Foundation (DFG) Emmy Noether Programme (CH 2676/1-1), the EU’s Horizon Europe project ARISE (Grant no.: 101135959), the German Federal Ministry of Education and Research (BMBF) project “RiG” (Grant no.: 16ME1001) and the European Research Council (ERC) project “SIREN” (Grant No.: 101163933). The authors gratefully acknowledge the scientific support and HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). The hardware is funded by the German Research Foundation (DFG).

References

- [1] R. Johansson and J. Flanagan. Tactile sensory control of object manipulation in humans. *The Senses: A Comprehensive Reference*, 6:67–86, 01 2010. doi:10.1016/B978-012370880-9.00346-7.
- [2] W. Yuan, S. Dong, and E. H. Adelson. GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force. *Sensors*, 17(12):2762, Dec. 2017. ISSN 1424-8220. doi:10.3390/s17122762. URL <https://www.mdpi.com/1424-8220/17/12/2762>.
- [3] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, et al. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845, 2020.
- [4] B. Ward-Cherrier, N. Pestell, L. Cramphorn, B. Winstone, M. E. Giannaccini, J. Rossiter, and N. F. Lepora. The tactip family: Soft optical tactile sensors with 3d-printed biomimetic morphologies. *Soft robotics*, 5(2):216–227, 2018.
- [5] S. Dong, D. K. Jha, D. Romeres, et al. Tactile-rl for insertion: Generalization to objects of unknown geometry. In *ICRA*, 2021.
- [6] R. Calandra, A. Owens, M. Upadhyaya, W. Yuan, J. Lin, E. H. Adelson, and S. Levine. The feeling of success: Does touch sensing help predict grasp outcomes?, 2025. URL <https://arxiv.org/abs/1710.05512>.
- [7] E. Helmut, N. Funk, T. Schneider, C. d. Farias, and J. Peters. Tactile-Conditioned Diffusion Policy for Force-Aware Robotic Manipulation, Oct. 2025. URL <http://arxiv.org/abs/2510.13324>. arXiv:2510.13324 [cs].
- [8] J. Hansen, F. Hogan, D. Rivkin, D. Meger, M. Jenkin, and G. Dudek. Visuotactile-RL: Learning Multimodal Manipulation Policies with Deep Reinforcement Learning. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8298–8304, May 2022. doi:10.1109/ICRA46639.2022.9812019. URL <https://ieeexplore.ieee.org/document/9812019>.
- [9] K. Yu, Y. Han, Q. Wang, V. Saxena, D. Xu, and Y. Zhao. MimicTouch: Leveraging Multimodal Human Tactile Demonstrations for Contact-rich Manipulation, Oct. 2023. URL <https://arxiv.org/abs/2310.16917v4>.
- [10] M. Yang, A. Church, Y. Lin, C. J. Ford, H. Li, E. Psomopoulou, D. A. Barton, N. F. Lepora, et al. Anyrotate: Gravity-invariant in-hand object rotation with sim-to-real touch. In *Conference on Robot Learning*, pages 4727–4747. PMLR, 2025.
- [11] B. Romero, H.-S. Fang, P. Agrawal, and E. Adelson. Eyesight hand: Design of a fully-actuated dexterous robot hand with integrated vision-based tactile sensors and compliant actuation. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1853–1860. IEEE, 2024.

- [12] N. Funk, E. Helmut, G. Chalvatzaki, R. Calandra, and J. Peters. Evetac: An Event-based Optical Tactile Sensor for Robotic Manipulation, Aug. 2024. URL <http://arxiv.org/abs/2312.01236>. arXiv:2312.01236 [cs].
- [13] D. Yin, S. Lu, J. Yang, Y. Zhang, Z. Dai, D. Nan, B. Cai, S. He, and X. Chen. Gelevent—a novel high-speed tactile sensor with event camera. *IEEE Transactions on Instrumentation and Measurement*, 74:1–13, 2025. doi:10.1109/TIM.2025.3551440.
- [14] 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*, 36(6):1619–1634, 2020.
- [15] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine. More than a feeling: Learning to grasp and regrasp using vision and touch. *IEEE Robotics and Automation Letters*, 3(4):3300–3307, 2018.
- [16] H. Qi, B. Yi, S. Suresh, M. Lambeta, Y. Ma, R. Calandra, and J. Malik. General in-hand object rotation with vision and touch. In *Conference on Robot Learning*, pages 2549–2564. PMLR, 2023.
- [17] B. Huang, Y. Wang, X. Yang, Y. Luo, and Y. Li. 3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing. In *Conference on Robot Learning*, pages 2557–2578. PMLR, 2025.
- [18] 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, Dec. 2022. URL <http://arxiv.org/abs/2212.03858>. arXiv:2212.03858 [cs].
- [19] T. Ablett, O. Limoyo, A. Sigal, A. Jilani, J. Kelly, K. Siddiqi, F. Hogan, and G. Dudek. Multimodal and force-matched imitation learning with a see-through visuotactile sensor. *IEEE Transactions on Robotics*, 41:946–959, 2024.
- [20] N. Funk, C. Chen, T. Schneider, G. Chalvatzaki, R. Calandra, and J. Peters. On the Importance of Tactile Sensing for Imitation Learning: A Case Study on Robotic Match Lighting. *IEEE Robotics and Automation Letters*, 11(5):6218–6225, 2026.
- [21] H. Xue, J. Ren, W. Chen, G. Zhang, Y. Fang, G. Gu, H. Xu, and C. Lu. Reactive Diffusion Policy: Slow-Fast Visual-Tactile Policy Learning for Contact-Rich Manipulation, Apr. 2025. URL <http://arxiv.org/abs/2503.02881>. arXiv:2503.02881 [cs].
- [22] N. Cheng, J. Xu, C. Guan, J. Gao, W. Wang, Y. Li, F. Meng, J. Zhou, B. Fang, and W. Han. Touch100k: A large-scale touch-language-vision dataset for touch-centric multimodal representation. *Information Fusion*, 124:103305, Dec. 2025. ISSN 1566-2535. doi:10.1016/j.inffus.2025.103305. URL <https://www.sciencedirect.com/science/article/pii/S1566253525003781>.
- [23] C. Higuera, A. Sharma, T. Fan, C. K. Bodduluri, B. Boots, M. Kaess, M. Lambeta, T. Wu, Z. Liu, F. R. Hogan, and M. Mukadam. Tactile Beyond Pixels: Multisensory Touch Representations for Robot Manipulation, June 2025. URL <http://arxiv.org/abs/2506.14754>. arXiv:2506.14754 [cs].
- [24] A. George, S. Gano, P. Katragadda, and A. B. Farimani. VITaL Pretraining: Visuo-Tactile Pretraining for Tactile and Non-Tactile Manipulation Policies, Sept. 2024. URL <http://arxiv.org/abs/2403.11898>. arXiv:2403.11898 [cs].
- [25] Z. Zhao, S. Haldar, J. Cui, L. Pinto, and R. Bhirangi. Touch begins where vision ends: Generalizable policies for contact-rich manipulation, June 2025. URL <http://arxiv.org/abs/2506.13762>. arXiv:2506.13762 [cs].

- [26] V. Dave, F. Lygerakis, and E. Rueckert. Multimodal Visual-Tactile Representation Learning through Self-Supervised Contrastive Pre-Training, Jan. 2024. URL <http://arxiv.org/abs/2401.12024>. arXiv:2401.12024 [cs].
- [27] L. Fu, H. Huang, L. Berscheid, H. Li, K. Goldberg, and S. Chitta. Safe Self-Supervised Learning in Real of Visuo-Tactile Feedback Policies for Industrial Insertion, Mar. 2023. URL <http://arxiv.org/abs/2210.01340>. arXiv:2210.01340 [cs].
- [28] C. Zhang, P. Hao, X. Cao, X. Hao, S. Cui, and S. Wang. VTLA: Vision-Tactile-Language-Action Model with Preference Learning for Insertion Manipulation, May 2025. URL <http://arxiv.org/abs/2505.09577>. arXiv:2505.09577 [cs].
- [29] G. Ye, Z. Zhang, X. Zhao, S. Wu, H. Lu, S. Lu, and H. Liu. Learning to Feel the Future: DreamTacVLA for Contact-Rich Manipulation, Dec. 2025. URL <http://arxiv.org/abs/2512.23864>. arXiv:2512.23864 [cs].
- [30] R. Feng, J. Hu, W. Xia, T. Gao, A. Shen, Y. Sun, B. Fang, and D. Hu. AnyTouch: Learning Unified Static-Dynamic Representation across Multiple Visuo-tactile Sensors, Apr. 2025. URL <http://arxiv.org/abs/2502.12191>. arXiv:2502.12191 [cs].
- [31] J. Zhao, Y. Ma, L. Wang, and E. H. Adelson. Transferable Tactile Transformers for Representation Learning Across Diverse Sensors and Tasks, Oct. 2024. URL <http://arxiv.org/abs/2406.13640>. arXiv:2406.13640 [cs].
- [32] 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. *CoRR*, abs/1810.10191, 2018. URL <http://arxiv.org/abs/1810.10191>. eprint: 1810.10191.
- [33] 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, July 2019. URL <http://arxiv.org/abs/1907.13098>. arXiv:1907.13098 [cs].
- [34] R. Feng, D. Hu, W. Ma, and X. Li. Play to the Score: Stage-Guided Dynamic Multi-Sensory Fusion for Robotic Manipulation, Oct. 2024. URL <http://arxiv.org/abs/2408.01366>. arXiv:2408.01366 [cs].
- [35] 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, Nov. 2023. URL <http://arxiv.org/abs/2311.00924>. arXiv:2311.00924 [cs].
- [36] M. Lambeta, T. Wu, A. Sengul, V. R. Most, N. Black, K. Sawyer, R. Mercado, H. Qi, A. Sohn, B. Taylor, N. Tydingco, G. Kammerer, D. Stroud, J. Khatha, K. Jenkins, K. Most, N. Stein, R. Chavira, T. Craven-Bartle, E. Sanchez, Y. Ding, J. Malik, and R. Calandra. Digitizing touch with an artificial multimodal fingertip, 2024. URL <https://arxiv.org/abs/2411.02479>.
- [37] J. Zhao, N. Kuppuswamy, S. Feng, B. Burchfiel, and E. Adelson. Polytouch: A robust multimodal tactile sensor for contact-rich manipulation using tactile-diffusion policies. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 104–110. IEEE, 2025.
- [38] P. Jin, B. Huang, W. W. Lee, T. Li, and W. Yang. Visual-Force-Tactile Fusion for Gentle Intricate Insertion Tasks. *IEEE Robotics and Automation Letters*, 9(5):4830–4837, May 2024. ISSN 2377-3766. doi:10.1109/LRA.2024.3379803. URL <https://ieeexplore.ieee.org/document/10476678/>.

- [39] K. Yu, Y. Han, Q. Wang, V. Saxena, D. Xu, and Y. Zhao. MimicTouch: Leveraging Multi-modal Human Tactile Demonstrations for Contact-rich Manipulation, Feb. 2025. URL <http://arxiv.org/abs/2310.16917>. arXiv:2310.16917 [cs].
- [40] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky. π_0 : A vision-language-action flow model for general robot control, 2026. URL <https://arxiv.org/abs/2410.24164>.
- [41] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow matching for generative modeling, 2023. URL <https://arxiv.org/abs/2210.02747>.
- [42] T. Xiao, M. Singh, E. Mintun, T. Darrell, P. Dollár, and R. Girshick. Early convolutions help transformers see better, 2021. URL <https://arxiv.org/abs/2106.14881>.
- [43] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need, 2023. URL <https://arxiv.org/abs/1706.03762>.
- [44] J. Urain, A. Mandlekar, Y. Du, N. Muhammad “Mahi” Shafiullah, D. Xu, K. Fragkiadaki, G. Chalvatzaki, and J. Peters. A survey on deep generative models for robot learning from multimodal demonstrations. *IEEE Transactions on Robotics*, 42:60–79, 2026. doi:10.1109/TRO.2025.3631816.
- [45] 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*, 44(10-11):1684–1704, 2025.
- [46] W. Peebles and S. Xie. Scalable diffusion models with transformers, 2023. URL <https://arxiv.org/abs/2212.09748>.
- [47] S. Dasari, O. Mees, S. Zhao, M. K. Srirama, and S. Levine. The ingredients for robotic diffusion transformers, 2024. URL <https://arxiv.org/abs/2410.10088>.
- [48] R. Krohn, V. Prasad, G. Tiboni, and G. Chalvatzaki. Self-supervised multisensory pretraining for contact-rich robot reinforcement learning. *IEEE Robotics and Automation Letters*, 11(6): 6799–6806, 2026. doi:10.1109/LRA.2026.3681156.
- [49] A. Nagrani, S. Yang, A. Arnab, A. Jansen, C. Schmid, and C. Sun. Attention bottlenecks for multimodal fusion, 2022. URL <https://arxiv.org/abs/2107.00135>.
- [50] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. URL <https://arxiv.org/abs/2010.11929>.
- [51] B. Tang, M. A. Lin, I. Akinola, A. Handa, G. S. Sukhatme, F. Ramos, D. Fox, and Y. Narang. Industreal: Transferring contact-rich assembly tasks from simulation to reality, 2023. URL <https://arxiv.org/abs/2305.17110>.
- [52] M. Heo, Y. Lee, D. Lee, and J. J. Lim. Furniturebench: Reproducible real-world benchmark for long-horizon complex manipulation. In *Robotics: Science and Systems*, 2023.
- [53] M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains, 2020. URL <https://arxiv.org/abs/2006.10739>.

A Appendix

This appendix provides supplementary material that supports the experimental evaluation presented in the main text.

A.1 Controller

We use a Cartesian pose controller implemented as a Cartesian-to-joint target generator coupled with a low-level joint impedance servo. At each control step (15Hz), the input, provided either by a SpaceMouse teleoperation interface or a learned policy, specifies a target end-effector pose in a reduced Cartesian space (x, y, z, ψ) , while roll and pitch are held fixed to values latched at reset. The Cartesian pose error is converted to a bounded Cartesian increment and mapped to joint motion using a Jacobian-based inverse kinematics module with joint-limit handling and nullspace regularization. This yields a desired 7-DoF joint configuration $\mathbf{q}_d \in \mathbb{R}^7$ that is tracked by a joint impedance controller running at 1kHz, applying $\tau = K_q(\mathbf{q}_d - \mathbf{q}) + K_{qd}(\dot{\mathbf{q}}_d - \dot{\mathbf{q}})$, resulting in compliant tracking of the commanded Cartesian targets.

Table 3: Controller parameters.

Parameter	Value
Outer-loop command rate	15 Hz
Command space	Cartesian position (xyzyaw)
Max linear increment per step	0.075 m
Max rotational increment per step	0.15 rad
Max joint increment per step	0.2 rad (per joint)
Nullspace gain	0.025
Regularization weight	1×10^{-2}
Max Cartesian velocity-IK iterations	300
Max nullspace iterations	300
Min. distance to joint limits	0.3 rad
Impedance servo rate	1000 Hz
Joint stiffness K_q	[40, 30, 50, 25, 35, 25, 10]
Joint damping K_{qd}	[4, 6, 5, 5, 3, 2, 1]

A.2 Parameter Count and Inference

We evaluate all nine policy variants on an NVIDIA GeForce RTX 4080, measuring inference speed over 100 runs per model and reporting the average runtime with 10 flow-matching integration steps (see Table 4). For each model we also report the number of trainable parameters, split into the perception encoder (Perc.) and the total model size. Every model is using the same policy head with 13,701,508 parameters.

Table 4: Parameter count and inference benchmark.

Model	Sensors	Co-Train	#params Perc.	#params Total	Tok/obs	ms	Hz
ViT-CNN	V	–	3,079,232	16,780,740	260	35.5	227.4
ViT	V	–	4,051,328	17,752,836	260	34.4	232.2
Sparsh	V+G	–	9,398,744	23,100,252	448	39.2	203.5
Sparsh	V+G	✓	15,500,056	29,201,564	448	38.9	208.4
Sparsh	V+G+E	–	15,500,056	29,201,564	640	54.6	153.4
MiTAS	V+E	–	8,769,280	22,470,788	452	35.6	223.7
MiTAS	V+G	–	5,074,776	18,776,284	452	35.3	224.1
MiTAS	V+G	✓	10,764,824	24,466,332	452	36.1	230.1
MiTAS	V+G+E	–	10,764,824	24,466,332	644	43.5	180.5

A.3 AdaLN-Zero time conditioning.

Each DiT block is modulated by the time embedding $\tau(t)$ using AdaLN-Zero [46, 47, 53]. For a sublayer input \mathbf{Y} , we apply adaptive layer normalization $\text{AdaLN}(\mathbf{Y}; \beta, \gamma) = \text{LN}(\mathbf{Y}) \odot (1 + \gamma) + \beta$, where (β, γ) are produced from $\tau(t)$. A per-block MLP outputs (β, γ, g) triplets for self-attention, cross-attention, and the MLP, and each residual update is gated as $\mathbf{Y} \leftarrow \mathbf{Y} + g f(\text{AdaLN}(\mathbf{Y}; \beta, \gamma))$. We initialize the AdaLN projection and the final output head to zero so the network starts near an identity mapping and learns time-conditioning smoothly.

A.4 Symbols and Hyperparameters

Table 5: Symbols of the MiTaS pipeline.

Symbol	Meaning
H	Action horizon
d_a	Action dimension
$\mathbf{A} \in \mathbb{R}^{H \times d_a}$	Action chunk
$\bar{\mathbf{A}} \in \mathbb{R}^{H \times d_a}$	Normalized action chunk
n_{act}	Number of executed actions per replanning step ($n_{\text{act}} \leq H$)
$O^{\text{vision}} \in \mathbb{R}^{2 \times 128 \times 128}$	Wrist-camera observation
$O^{\text{gelsight}} \in \mathbb{R}^{2 \times 120 \times 160}$	GelSight tactile observation
$O^{\text{evetac}} \in \mathbb{R}^{16 \times 120 \times 160}$	EveTac observation
$\mathbf{x}_0 \in \mathbb{R}^{H \times d_a}$	Initial noise sample, $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
$\mathbf{x}_1 \in \mathbb{R}^{H \times d_a}$	Data endpoint for flow matching ($\mathbf{x}_1 = \bar{\mathbf{a}}$)
t	Continuous flow time, $t \in [0, 1]$
$\mathbf{x}_t \in \mathbb{R}^{H \times d_a}$	On-path state returned by the conditional flow matcher
$\mathbf{u}_t \in \mathbb{R}^{H \times d_a}$	Target conditional velocity returned by the flow matcher
$\mathbf{C} \in \mathbb{R}^{N_{\text{cond}} \times D}$	Conditioning tokens from the multisensory encoder
N_{cond}	Number of conditioning tokens
D	Token embedding width (shared by \mathbf{C} and action tokens)
$\mathbf{z}_i \in \mathbb{R}^D$	Sensor token
$\mathbf{p}_{g(i)} \in \mathbb{R}^D$	Learned position embedding for token-grid location $g(i)$
$\mathbf{e}_{s(i)} \in \mathbb{R}^D$	Learned modality embedding for source modality $s(i)$
$s(i) \in \{\text{vision}, \text{gelsight}, \text{evetac}\}$	Modality identifier for token i
$\mathbf{z}'_i \in \mathbb{R}^D$	Augmented token, $\mathbf{z}'_i = \mathbf{z}_i + \mathbf{p}_{g(i)} + \mathbf{e}_{s(i)}$
$\boldsymbol{\tau}(t) \in \mathbb{R}^{d_t}$	Time embedding of t
$v_\theta(\mathbf{x}_t, t, \mathbf{C})$	DiT network predicting velocity / vector field
$\hat{\mathbf{v}}_\theta \in \mathbb{R}^{H \times d_a}$	Predicted velocity chunk
m_h	Validity mask for timestep h (episode padding mask)
\mathcal{L}_{FM}	Flow-matching training loss (masked MSE)
K	Number of Euler integration steps at inference
Δt	Euler step size, $\Delta t = 1/K$
$\Delta \mathbf{p} \in \mathbb{R}^4$	Delta position and delta yaw: $[\Delta x, \Delta y, \Delta z, \Delta \psi]$

Table 6: Hyperparameters.

Hyperparameter	Value	Hyperparameter	Value
Batch size	64	Vision normalization	$[0, 1]$
Action horizon H	16	Gelsight normalization	$[0, 1]$
Executed steps per replan n_{act}	1-3	Evetac normalization	$[-\frac{1}{2}, \frac{1}{2}]$
Observation window n_{obs}	4	Embedding width D	256
Action normalization	minmax	DiT depth L	10
Action target	$\Delta \mathbf{p} \in \mathbb{R}^4$	Attention heads	8
Time embedding dim d_t	128	MLP ratio	4
Time embedding	sinusoidal	Dropout	0.0
Flow-matching noise σ	0.0	FM inference steps K	10
Policy learning rate	1×10^{-4}	Encoder learning rate	1×10^{-5}
Weight decay	1×10^{-6}	Adam betas	(0.95, 0.999)
Learning-rate schedule	cosine	Gradient clipping	1.0

A.5 Task Progress

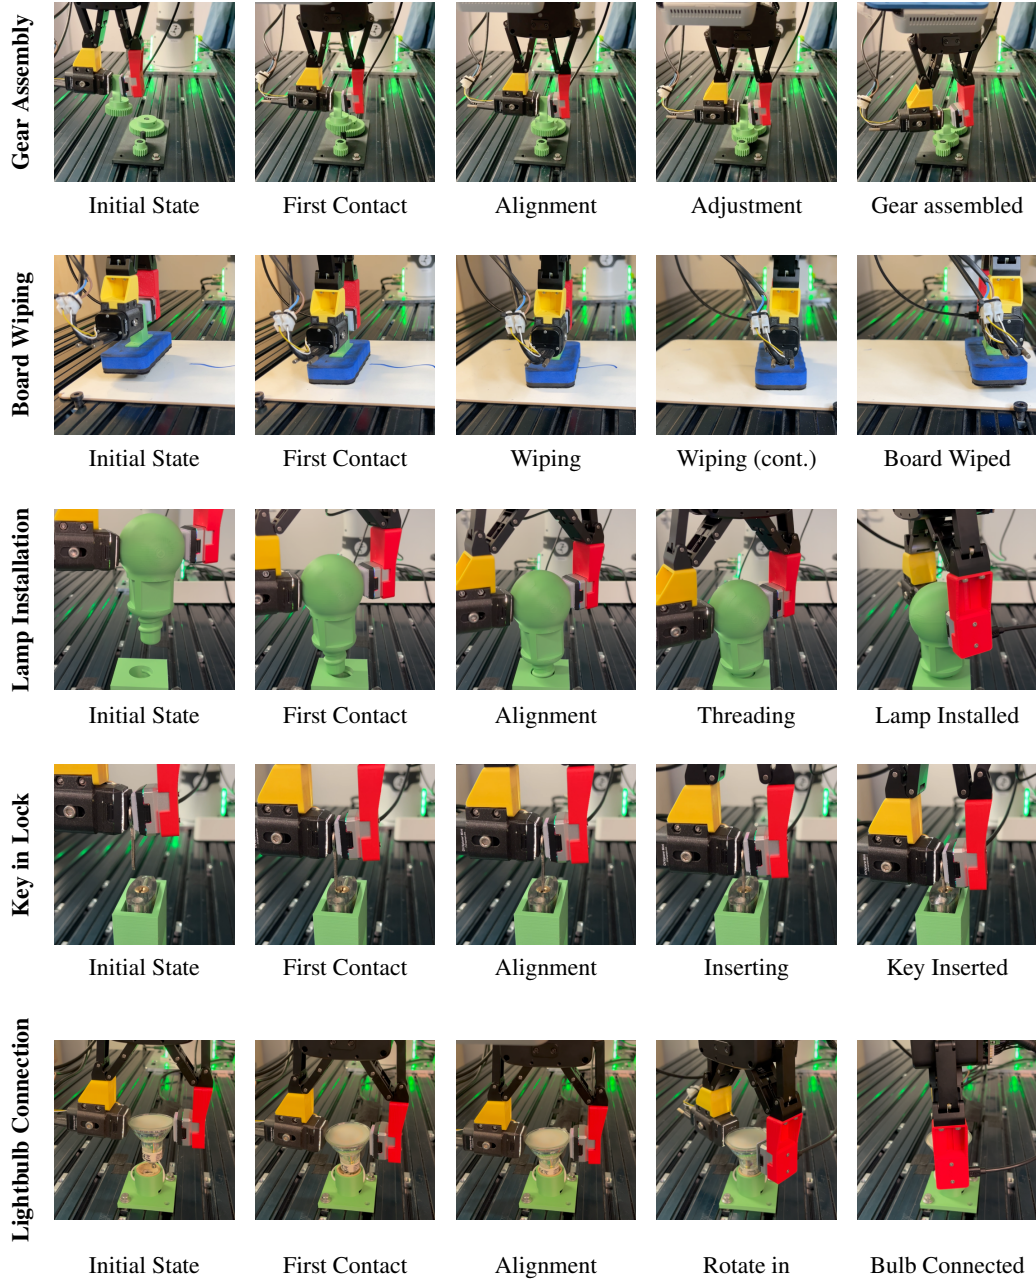


Figure 6: Task progress across five manipulation tasks. Each row shows representative frames from a successful rollout, progressing from the initial scene configuration (left) to task completion (right). The tasks span a range of contact-rich skills, including Gear Assembly, Board Wiping, Lamp Installation (FurnitureBench), Key in Lock, and Lightbulb Connection.

A.6 Failure Cases

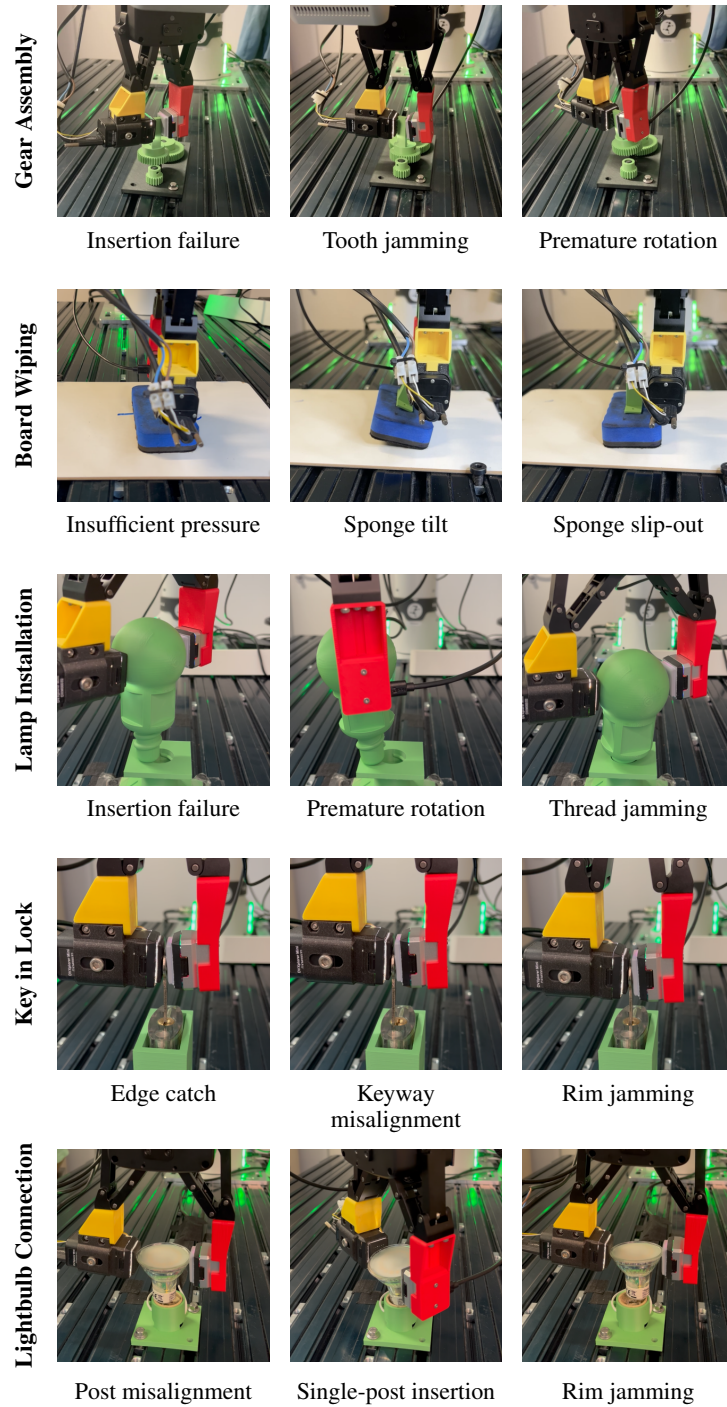


Figure 7: Examples of failure cases across the five manipulation tasks. Each row shows representative frames from a failed rollout. These examples are illustrative and not exhaustive.

A.7 In-Hand Distribution of Initial States

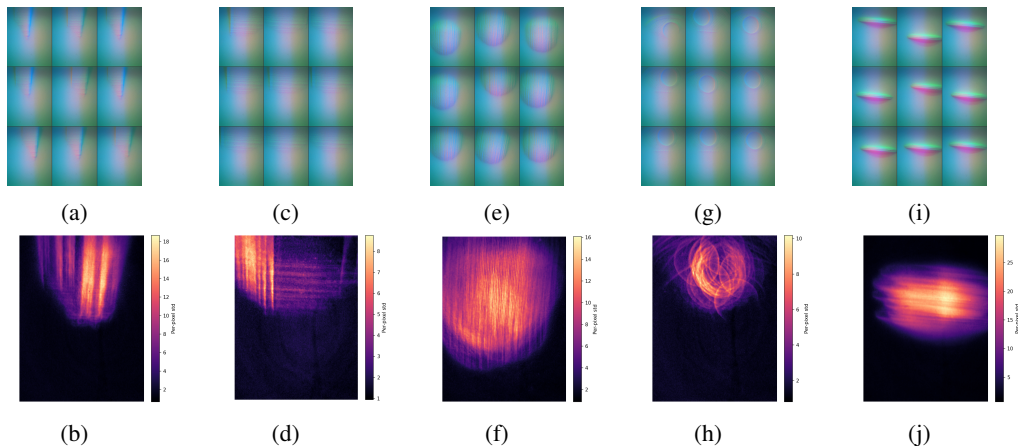


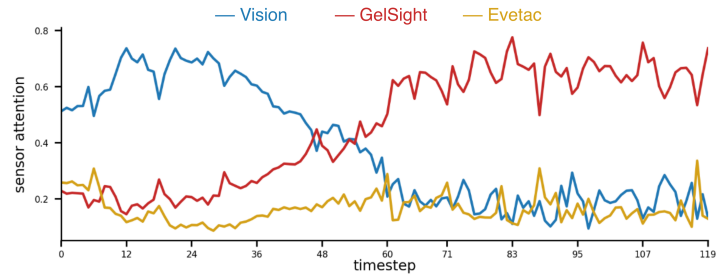
Figure 8: In-hand GelSight distributions after reset across Gear Assembly (a-b), Board Wiping (c-f), Lamp Install (e-f), Key In Lock (g-h) and Lightbulb Connect (i-j). Each column shows a 3×3 montage on top and the per-pixel standard deviation at the bottom of the respective training dataset.

A.8 Results overview

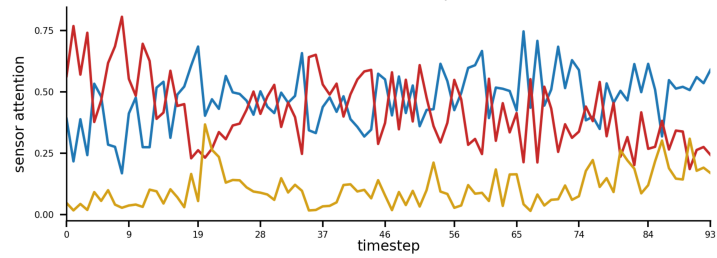
Table 7: Policy performance across all manipulation tasks (success rate in %). Bold: best in column. Avg.: mean over the five evaluated tasks.

Model	Sensors	Co-Train	Gear Assembly	Board Wiping	Lamp installation	Key in Lock	Lightbulb Connection	Avg.
ViT-CNN	V	–	40%	70%	20%	0%	0%	26%
ViT	V	–	40%	60%	55%	0%	0%	31%
MiTAS	V+E	–	0%	80%	65%	0%	40%	37%
Sparsh	V+G	–	25%	50%	55%	20%	15%	33%
MiTAS	V+G	–	45%	70%	65%	55%	20%	51%
Sparsh	V+G	✓	15%	85%	70%	5%	10%	37%
MiTAS	V+G	✓	55%	85%	60%	20%	45%	53%
Sparsh	V+G+E	–	50%	85%	70%	10%	55%	54%
MiTAS (Ours)	V+G+E	–	90%	90%	80%	75%	65%	80%

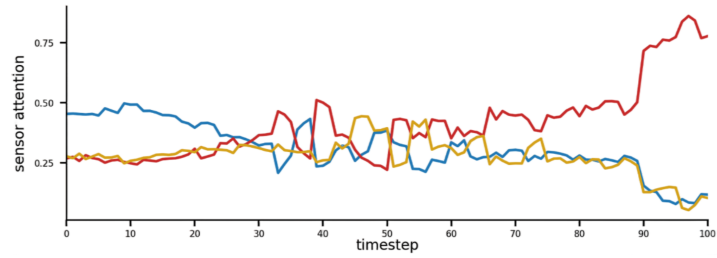
A.9 Attention analysis



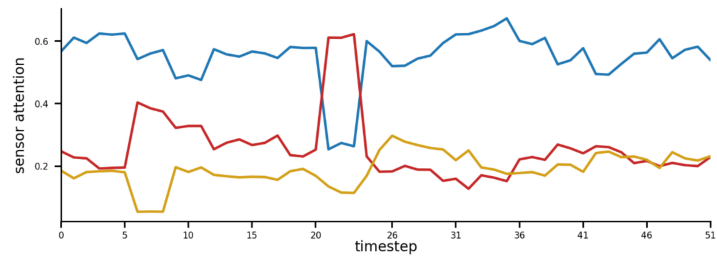
(a) Gear Assembly



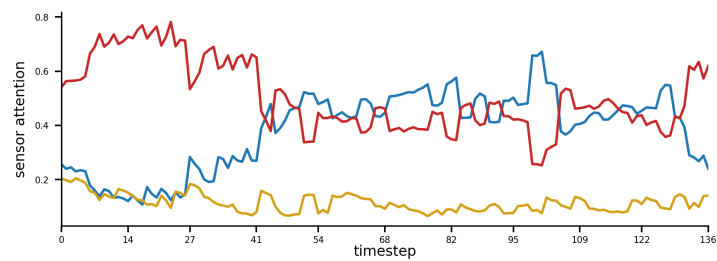
(b) Board Wiping



(c) Lamp Installation




(d) Key in Lock




(e) Lightbulb Connection

Figure 9: For each task we plot the cross-attention from the executed action token to each sensor modality (Vision, GelSight, Evetac) over the course of a successful episode.


A.10 Evetac Activation




(a) Gear Assembly




(b) Board Wiping



(c) Lamp Installation



(d) Key in Lock



(e) Lightbulb Connection

Figure 10: We define the Evetac activation as the L1-distance to the base image in pixel-space. Evetac activations differ for each task and task-stage indicating its usefulness to the policy. During reaching, beginning of a trajectory, the activation is low, it spikes by contact and provides high frequency features during object interaction.