

TacUMI: A Multi-Modal Universal Manipulation Interface for Contact-Rich Tasks

Tailai Cheng^{1,2}, Kejia Chen¹, Lingyun Chen¹, Liding Zhang¹, Yue Zhang¹, Yao Ling¹,
Mahdi Hamad², Zhenshan Bing^{3,1}, Fan Wu^{4,1}, Karan Sharma², Alois Knoll¹

Abstract—Task decomposition is critical for understanding and learning complex long-horizon manipulation tasks. Especially for tasks involving rich physical interactions, relying solely on visual observations and robot proprioceptive information often fails to reveal the underlying event transitions. This raises the requirement for efficient collection of high-quality multi-modal data as well as robust segmentation method to decompose demonstrations into meaningful modules.

Building on the idea of the handheld demonstration device Universal Manipulation Interface (UMI), we introduce TacUMI, a multi-modal data collection system that integrates additionally ViTac sensors, force–torque sensor, and pose tracker into a compact, robot-compatible gripper design, which enables synchronized acquisition of all these modalities during human demonstrations. We then propose a multi-modal segmentation framework that leverages temporal models to detect semantically meaningful event boundaries in sequential manipulations. Evaluation on a challenging cable mounting task shows more than 90% segmentation accuracy and highlights a remarkable improvement with more modalities, which validates that TacUMI establishes a practical foundation for both scalable collection and segmentation of multi-modal demonstrations in contact-rich tasks. The design and experiment results are available at our project website: <https://tac-umi.github.io/TacUMI/>.

I. INTRODUCTION

In real-world applications such as industrial assembly and household services, robots are often required to perform long-horizon tasks consisting of a sequence of modular skills. However, learning such long-horizon tasks, especially when there are intensive physical interactions, remains a fundamental challenge. While recent imitation learning frameworks, such as ACT [1] and Diffusion Policy [2], have achieved strong performance on short-horizon tasks, they treat the whole demonstration as a monolithic sequence without explicit segmentation, which makes them struggle to learn long-horizon tasks at once [3]. In contrast, learning separately modular motor skills and high-level coordinator to compose them have proven to be an efficient and promising paradigm [4], [5].

This hierarchical learning paradigm requires firstly decomposing a long demonstration into semantically meaningful, modular skills. In this segmentation, tactile information can play an important role, as it can capture contact dynamics that may be missed by visual observation [6]. For example,

¹ School of Computation, Information and Technology, Technical University of Munich, Germany.

² Agile Robots SE, Munich, Germany

³ State Key Laboratory for Novel Software Technology and the School of Science and Technology, Nanjing University (Suzhou Campus), China.

⁴ Shanghai University, China.

Corresponding author: Kejia Chen (kejia.chen@tum.de) and Zhenshan Bing (zhenshan.bing@tum.de).




Fig. 1. Overall setup for the cable mounting task. The left hand holds a modified UMI gripper equipped with a ViTac sensor and a Vive Tracker, while the right hand holds the proposed TacUMI gripper for multimodal data collection. A camera positioned in the center records third-person visual information during the demonstrations.

during cable mounting [7], operators often stretch the cable to create tension. This step is difficult to discern visually but can be clearly captured through tactile signals. However, tactile sensing has not been sufficiently explored in existing research due to limitation in both data acquisition and learning methods.

In terms of data collection, teleoperation-based systems [8]–[10] provide precise demonstrations, but their high cost, lack of intuitive control, and limited force feedback severely hinder practical usability. Recent handheld data collection devices, such as the Universal Manipulation Interface (UMI) [11] and FastUMI [12], have lowered the barrier for human-in-the-loop demonstration. Nevertheless, they rely solely on vision and pose, and the critical tactile information is missing from their design. When it comes to learning methods, visual and proprioceptive data have been widely explored for task decomposition [4], [13], [14]. However, less attention were paid to the utilization and integration of tactile sensing.

To address these limitations, we present a new solution that tightly integrates multi-modal sensing hardware with a task segmentation learning framework. On the hardware side, we develop an handheld data collection device inspired by UMI, but enhanced with ViTac sensors on fingertips, a 6D force-torque (F/T) sensor on the wrist, and a high-precision 6D pose tracker. Especially for F/T signals, we introduce a continuously lockable jaw mechanism that eliminates the user's interference when grasping. This enables synchronize collecting clean measurement of contact tactile, forces/torque, and gripper pose information during demonstrations.

On the algorithmic side, we propose a BiLSTM-based task segmentation framework that learns to partition long-horizon demonstrations into a sequence of skills. Using the rich temporal structure and complementary signals from multiple modalities, our approach identifies transition boundaries between each stage of long-horizon manipulation tasks. This segmentation lays a foundation for learning modular skills and composing them, which makes learning from long-horizon tasks more scalable and interpretable.

To validate our proposed data collection and segmentation approach, we evaluate our system on a challenging cable mounting task. In comparison to other UMI-based designs as well as teleoperation, TacUMI obtains uninterfered F/T data in the most efficient way. By integrating multiple, especially the tactile, modalities collected by TacUMI, the learned segmentation model reveals clear and semantically meaningful transition across different phases of the manipulation process. Moreover, we deploy the model trained from TacUMI-collected dataset on another dataset collected by teleoperating robots, and achieves comparable segmentation accuracy. These evaluations demonstrate the effectiveness of our system for interpreting and segmenting long-horizon manipulation tasks.

II. RELATED WORK

A. Data Collection Interfaces for Robot

High-quality multimodal data is crucial for understanding complex, contact-rich manipulation tasks. Conventional teleoperation systems [2], [6], [15] offer precise control and direct mapping to robot embodiments, but they are often costly, cumbersome to set up, and lack intuitive force feedback for the operator. These limitations have motivated the development of *handheld robotic data collection devices* that allow operators to physically manipulate objects in a natural manner while capturing sensor data.

The Universal Manipulation Interface (UMI) [11] and FastUMI [12] provide low-cost handheld tools that enable rapid collection of visual and pose data from human demonstrations. While effective for efficient data acquisition, both systems rely solely on vision-based sensing and thus cannot capture tactile feedback or the interaction forces between the gripper and the environment. Moreover, their pose estimation, built on SLAM-based methods, inevitably suffers from cumulative drift during long-horizon demonstrations. ForceMimic [16] extends data collection by incorporating a 6D force-torque (F/T) sensor and a ratchet-based locking mechanism, allowing the recording of wrench data during manipulation. However, its soft gripper structure and discrete locking mechanism limit the achievable clamping force and stability. As a result, the gripper can fail to provide sufficient normal force and friction at the contact interface, leading to slippage or even loss of the manipulated object under external loads.

In contrast, our device integrates synchronized *tactile sensing, force-torque measurement, and precise 6-DoF pose tracking* via a pose tracker, eliminating cumulative drift issues in long tasks. The continuous locking mechanism allows the operator to maintain a stable grasp without sustaining

trigger pressure, removing internal actuation forces from FT recordings and ensuring clean interaction data. Furthermore, our gripper’s high clamping force capability prevents slippage during high-tension interactions (e.g., deformable linear object(DLO) manipulation), extending applicability to a wider range of contact-rich scenarios.

B. Event Segmentation

Event segmentation aims to decompose continuous demonstrations into semantically meaningful skills, which is essential for learning and executing long-horizon tasks. Unsupervised methods such as Bottom-Up Skill Discovery [4] employ agglomerative clustering to build hierarchical task structures from raw demonstrations, identifying recurring motion patterns as candidate skills. XSkill [13] learns a shared skill representation space and a set of learnable skill prototypes from unlabelled human and robot videos, aligning them into a unified latent space for cross-embodiment skill discovery. SkIDRAW [14] jointly performs unsupervised skill discovery and trajectory segmentation directly from raw, unlabelled trajectories. While these approaches are effective for discovering latent skills without annotations, they are unsuitable for our cable mounting task, where each skill is explicitly defined and unsupervised clustering may lead to ambiguous or inconsistent segmentation.

Other segmentation works, such as [17], [18], operate by detecting motion boundaries from human or robot demonstrations, relying primarily on camera video and proprioceptive data. However, these modalities alone are insufficient for our scenario: in cable mounting, certain skills, such as *tensioning the cable*, involve minimal observable motion and thus cannot be reliably detected using only visual and proprioceptive cues.

Our segmentation approach follows the *object-centric* perspective from LEMMo-Plan [6], where segmentation is driven by changes in the state of the manipulated object rather than solely by end-effector kinematics. Although LEMMo-Plan showed that ViTac images can directly capture subtle skill transitions during object manipulation, we found that using tactile data alone was insufficient to achieve robust performance. To address this, we integrate ViTac data, force-torque measurements, and third-person visual observations, and process them using a BiLSTM-based temporal model. This multi-modal integration enables reliable detection of subtle contact-state changes and accurate segmentation of complex, long-horizon cable mounting demonstrations.

III. METHODOLOGY

A. Mechanical Design for Robot-Executable Demonstrations

We propose a new handheld demonstration device that supports multi-modal, synchronized, and robot-compatible data collection for complex manipulation tasks. Inspired by previous designs [11], [12], [16], our system significantly extends hardware capabilities for demonstrating long-horizon and contact-rich manipulation tasks. The mechanism preserves a robot wrist-end effector layout and integrates tactile sensing, 6D force-torque (F/T) sensing, and 6D pose tracking in a compact, one-handed form factor (Fig. 2)




Fig. 2. TacUMI structure. Exploded view of the proposed multimodal data collection gripper, showing: (1) ViTac sensor (GelSight Mini); (2) Vive tracker for 6-DoF pose tracking; (3) Trigger connected to the rack-and-pinion mechanism for grasp actuation; (4) Continuous locking part that enables the trigger to be locked at arbitrary jaw opening widths; (5) 6-axis Force/Torque (F/T) sensor; (6) Handle ergonomically designed for single-handed operation. Insets illustrate the locking mechanism in the locked and released states.

1) *Robot-Compatible Force–Torque Stack*: A 6-axis F/T sensor is mounted between the gripper and a rear handle shaped as a *standard robot flange* (Fig. 2 ⑤). This preserves the physical stack-up used on robots, so the measured wrench during human operation matches robot–environment interactions at execution. Because the F/T sensor sits between the wrist-equivalent interface and the tool, the recorded wrench is directly transferable to downstream control without frame re-ordering or ad-hoc calibration.

2) *Rack-and-Pinion Trigger with Continuous Self-Locking*: We retain the UMI-style rack-and-pinion trigger for intuitive force scaling, where finger closure is proportional to trigger effort [11]. In prior designs, continuous trigger pressing was required to maintain a grasp, and this actuation leaked into F/T measurements. To avoid this, we introduce a **continuous self-locking mechanism** (Fig. 2 ④) that rigidly fixes the jaw aperture at any opening. Once locked, the fingers and grasped object form a stiff extension of the tool; the F/T sensor then measures only *external* contact wrenches. Unlike ratcheted locks and soft fingertips in previous work [16], the continuous lock avoids quantization and enables high gripping forces across object sizes, maintaining hold even under significant tensile loads. Our software optionally filters any residual trigger-coupled components when only external contacts are desired for segmentation.

3) *ViTac Sensing for Contact-Rich Transitions*: Custom fingertip slots accept ViTac modules (Fig. 2 ①), which capture fine surface deformations (Fig. 5). These tactile cues disambiguate visually similar states and thus improve phase understanding and segmentation in contact-centric tasks.

4) *Drift-Free 6-DoF Pose Tracking*: A rigidly mounted Vive Tracker on the top cover provides accurate 6-DoF trajectories with negligible long-horizon drift, enabling a fixed transform to the tool center point (TCP). Compared with visual–inertial SLAM pipelines [11], this eliminates cumulative drift during extended demonstrations and yields stable trajectories for retargeting.

In summary, with a robot-compatible F/T stack with a *continuous* mechanical lock, the device isolates *external* interaction wrenches during demonstrations, which prior UMI-family designs lack. Integrated further with tactile and pose sensing, it is a compact, use-friendly and robot-compatible data collection tool which supports synchronized multi-modal sensing and is especially suitable for capturing contact-rich interactions in long-horizon demonstrations. Ergonomically, it supports single-handed operation: the center of mass sits near the palm, the trigger is indexed to the forefingers, and the lock is thumb-engaged with index release. Meanwhile, it preserves wrist–tool configuration which allows seamless transfer to real robots. The overall structure is lightweight and largely composed of 3D-printed components, making it easy to manufacture, customize, and reproduce.

B. Event Segmentation

We aim to segment long-horizon manipulation demonstrations temporally into meaningful short-skill segments using rich multi-modal sensory feedback. Our approach leverages synchronized *tactile images*, *force-torque data*, *third-person view RGB images*, and *pose data* to detect skill transitions with high granularity.

1) *Network Architecture*: At each timestep, we fuse representation from tactile, RGB image, F/T and pose (Fig. 3(a)). For tactile, we feed raw ViTac frames to a ResNet-50 backbone pretrained on ImageNet, truncate the classification head, and add a linear layer to yield a 256-D embedding. For the third-person views, inspired by Diffusion Policy [2], we use a ResNet-18 with all BatchNorm replaced by GroupNorm, remove global average pooling and fully connected layers, apply a Spatial softmax to extract spatial attention features, and project to a 256-D embedding. The raw 6-D F/T stream is preprocessed to remove operator trigger artifacts (*pull*, *lock*, *release*) using a secondary segmentation model (see Section III-C), yielding an interaction-only signal (Fig. 4b). The pose data obtained from trackers are transformed into the grippers’ tool center point (TCP). For each hand, we represent the TCP pose as a 7-D vector. Concatenating both left and right hand TCP poses yields a 14-D pose feature. The resulting vectors are normalized and concatenated with the tactile and visual embeddings to form

$$\mathbf{x}_t = [\mathbf{I}_{\text{tactile}}^{(t)} \parallel \mathbf{I}_{\text{3rd-view}}^{(t)} \parallel \mathbf{f}^{(t)} \parallel \mathbf{P}^{(t)}] \in \mathbb{R}^{532}, \quad (1)$$

where $\mathbf{I}_{(\cdot)}^{(t)} \in \mathbb{R}^{256}$, $\mathbf{f}^{(t)} \in \mathbb{R}^6$, and $\mathbf{P}^{(t)} \in \mathbb{R}^{14}$.

To capture temporal dependencies across multimodal sequences, we evaluate three different backbone architectures:

- a three-layer Bidirectional LSTM (BiLSTM) with hidden size 128 per direction. Taking the fused multi-modal representation \mathbf{x}_t at each timestep as input, the BiLSTM produces a sequence of hidden states $\{\mathbf{h}_t\}_{t=1}^T$, which are subsequently passed through a dropout layer and a linear classifier to produce $\mathbf{z}_t \in \mathbb{R}^C$, where C denotes the number of skills.
- a Temporal Convolutional Network (TCN) composed of stacked residual blocks with dilated 1D convolutions, GELU activations, and Layer Normalization. This

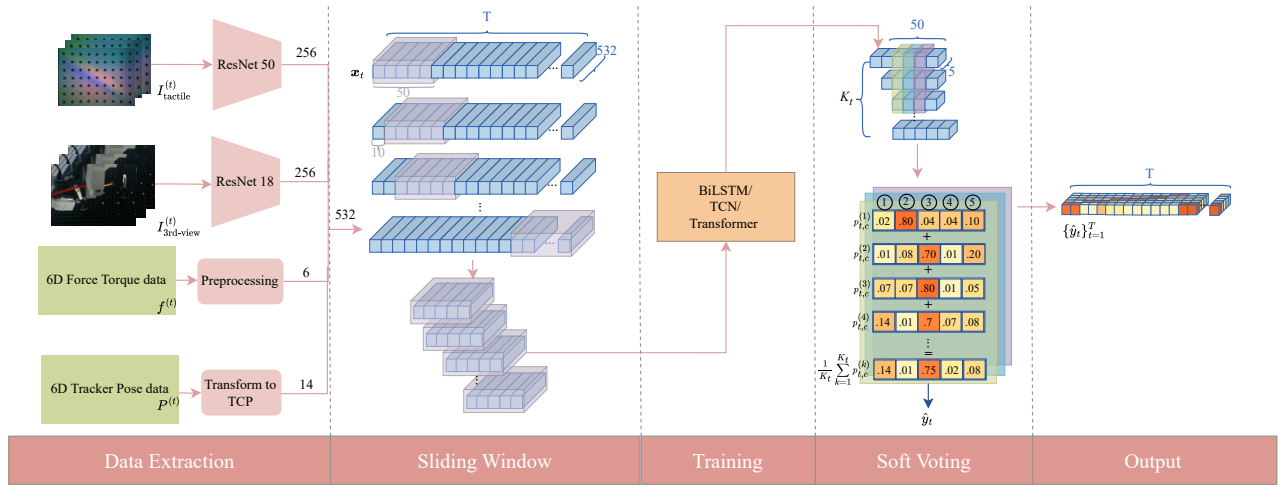


Fig. 3. **Pipeline of the proposed event segmentation algorithm.** The framework consists of four main stages: (a) **Data Extraction:** tactile images are processed by a ResNet50 to obtain 256-dimensional embeddings; third-person view images are processed by a ResNet18 with GroupNorm to obtain another 256-dimensional visual embedding; raw 6D F/T signals are preprocessed and 6D tracker pose are transformed to TCP. (b) **Sliding Window:** fused feature sequences of size $[T, 532]$ are segmented into overlapping windows of length 50 with stride 10. (c) **Training:** each window is fed into one of three sequence models (BiLSTM, TCN, or Transformer) to capture temporal dependencies and produce per-frame skill predictions. (d) **soft voting inference:** overlapping predictions for the same frame are aggregated by averaging class probabilities and assigning the highest-probability label, restoring the original sequence length for final output.

architecture enlarges the receptive field and captures multiscale temporal patterns while maintaining efficient computation.

- a Transformer Encoder with learned linear input projection, sinusoidal positional encoding, and L stacked encoder layers, each consisting of multi-head self-attention and feedforward sublayers.

All three architectures output frame-wise logits from sequence-to-sequence mapping $f: \mathbb{R}^{T \times 532} \rightarrow \mathbb{R}^{T \times C}$.

2) *Sliding Window and Training Strategy:* In sequential tasks, there is usually a label imbalance in terms of number of frames. Particularly, the *idle* class where nothing happens are often over-represented. To handle this imbalance and also to increase training efficiency, we adopt a sliding window [17] to extract fixed-length overlapping subsequences from each demonstration (Fig. 3(b)). This strategy both augments the dataset and improves the model’s focus on meaningful skill transitions.

We use a window size of 50 and a stride of 10. To ensure that the model learns from active manipulation segments, we discard any window in which more than 80% of the frames are labeled as *idle*, enforcing a minimum action ratio of 20%. Each retained window is paired with its corresponding frame-wise label sequence and used to supervise the temporal model. Training is conducted on the long-horizon task using the Adam optimizer with standard learning rate scheduling. To prevent overfitting, we apply dropout after the BiLSTM layers and employ early stopping based on validation segmentation accuracy.

3) *Soft Voting for Inference:* After the BiLSTM model predicts frame-wise skill labels for each input window, we obtain multiple overlapping predictions for the same frame due to the use of a sliding window during both training and inference. Each frame in the original demonstration may appear in several windows, each providing an independent classification output.

To reconcile these overlapping predictions and restore the temporal alignment with the original sequence, we apply a soft voting strategy. Specifically, for each frame t in the original sequence, we collect all predicted class probability vectors from the windows that include frame t . We aggregate probabilities from all overlapping windows that include t . Specifically, let K_t denote the number of such windows and $p_{t,c}^{(k)}$ the predicted probability of class c from the k -th window. We compute the averaged class probability and assign the final label by

$$\hat{y}_t = \arg \max_{c \in \{1, \dots, C\}} \frac{1}{K_t} \sum_{k=1}^{K_t} p_{t,c}^{(k)}, \quad (2)$$

so that \hat{y}_t corresponds to the class with the highest mean probability across windows. As shown in (2).

This soft aggregation method ensures temporal consistency, reduces prediction noise, and produces a smooth, frame-aligned label sequence for the entire demonstration. The final output preserves the original sequence length and provides a single semantic label per frame, suitable for downstream evaluation and visualization.

C. F/T Data Preprocessing

In this section, we describe two preprocessing steps applied to raw force-torque (F/T) signals before they are used for segmentation. Prior to these two steps, we apply Bota Systems’ inertial wrench compensation to remove payload gravity/inertia and sensor bias, yielding contact-only wrenches.

1) *Trigger Action Filtering:* A unique downstream use case of our segmentation model is to **filter out trigger-related transitions** in force-torque signals. As shown in Fig. 4(a), with help of the continuous locking mechanism, the collected force and torque signals drop sharply to near zero right after locking, which indicates that internal actuation effects have been eliminated and resembles the measurement




Fig. 4. Comparison of original and processed Force/Torque (F/T) data. (a) **Original F/T data:** The recorded signals reveal distinct operator-induced events, namely pulling and locking the trigger (blue-shaded region) followed by releasing it (green-shaded region), both of which cause significant disturbances in the measured force and torque signals. (b) **Processed F/T data:** The influence of these three actions has been removed, yielding clean signals that can be directly utilized by the robot without the undesired human-induced artifacts.

of wrist-mounted F/T sensors on robots. Nevertheless, there are still temporary interferences when the operator pulls the trigger, locks it or releases it. To filter these interferences out, we train a second instance of the same segmentation architecture to explicitly classify these intervals. We then replace the identified interference intervals with Gaussian noise whose mean and standard deviation estimated from the first 20 frames of each sequence, where no trigger actions occur. This noise injection preserves the natural statistical characteristics of the signal, as shown in Fig. 4(b), while effectively removing internal actuation effects. The resulting interaction-only force-torque stream maintains temporal continuity and is better suited for downstream robotic learning and control.

2) *Gripper-to-Robot F/T Data Mapping:* To adapt the gripper-collected F/T data for use across different robot platforms, we map the measured wrench from the gripper F/T sensor frame $\{G\}$ to the robot end-effector frame $\{R\}$. The mapping is given by

$$\mathbf{F}_R = R_{GR}\mathbf{F}_G, \quad (3)$$

$$\boldsymbol{\tau}_R = R_{GR}\boldsymbol{\tau}_G + \mathbf{r}_{GR} \times (R_{GR}\mathbf{F}_G), \quad (4)$$

where $\mathbf{F}_G \in \mathbb{R}^3$ and $\boldsymbol{\tau}_G \in \mathbb{R}^3$ denote the force and torque measured in $\{G\}$, $\mathbf{F}_R \in \mathbb{R}^3$ and $\boldsymbol{\tau}_R \in \mathbb{R}^3$ are the corresponding quantities expressed in $\{R\}$, $R_{GR} \in SO(3)$ is the rotation matrix (expressed in the TCP coordinate system) that transforms vectors from the gripper-mounted F/T sensor frame to the robot-mounted F/T sensor frame, and $\mathbf{r}_{GR} \in \mathbb{R}^3$ is the displacement vector (expressed in the TCP coordinate system) from the measurement center O_G of the gripper-mounted F/T sensor to the measurement center O_R of the robot-mounted F/T sensor.

Equations (3) and (4) follow the standard wrench transformation under a change of reference frame, and ensure that the interaction forces measured at the gripper are consistently expressed in the robot end-effector frame. This unified representation enables seamless transfer of the gripper-collected F/T data across robot platforms with varying wrist sensor

and TCP configurations.

IV. EXPERIMENTS

In this section, we perform a cable mounting experiment to validate both the correctness and efficiency of multimodal data collection using our proposed handheld gripper, and to demonstrate the effectiveness of the event segmentation algorithm.

A. Experimental Setup

1) *Hardware and Sensors:* We conduct experiments using our custom-designed handheld tools, whose multi-modal data collection system consists of four sensors: a Bota Systems SensONE 6D force-torque sensor, a GelSight Mini ViTac sensor, an HTC Vive Tracker for 6-DoF pose tracking, and a fixed-position RealSense D435i camera for third-person visual observation. These sensors operate at different frequencies: the force-torque sensor at 1000 Hz, the ViTac sensor (with online processing) at 16.67 Hz, the Vive Tracker at 60 Hz, and the RGB-D camera also at 60 Hz. To ensure temporal alignment across modalities, we down-sample all sensor streams to a common frequency of 16.67 Hz, matching that of the tactile sensor, the lowest among all devices.

For dual-arm manipulation tasks, we distribute different sensing modalities to different hands. For example, in the cable mounting task (shown in Fig. 5), the right-hand TacUMI is equipped with a force-torque sensor and a Vive Tracker to maintain stable grasping and guide the cable along its path, while the left-hand gripper carries a ViTac sensor and a Vive Tracker to perform fine manipulations, including tensioning, positioning, and inserting the cable into clips. During each trial, the operator progresses along the clip sequence, coordinating both grippers to maintain cable stability while executing precise insertions. Variations in clip position and orientation introduce non-repetitive transitions between subtasks, providing a challenging benchmark for multimodal task segmentation and validating the robustness of our proposed approach.

To validate that the multimodal demonstrations collected with our handheld gripper can be effectively transferred to a robot platform, we additionally collect a test dataset from teleoperation, where we employ *sigma.7* haptic devices to control two *Franka Emika* robots in a master-slave configuration. The teleoperation setup mirrors the handheld gripper setup: the guiding arm is equipped with a wrist-mounted F/T sensor, and the manipulating arm is equipped with a fingertip-mounted tactile sensor, so it captures the same modalities (tactile, force-torque, and pose from proprioceptive data) as the handheld device. We then evaluate on the teleoperation dataset whether the segmentation models trained exclusively on handheld-gripper data could generalize well to robot-collected data.

B. Data Collection Correctness and Efficiency

To evaluate the correctness of F/T data collected by different gripper designs, we compared four variants as illustrated in Fig. 6. The first design (Fig. 6(a)) places the F/T sensor between the handle and the gripper body, but does not include any locking mechanism, so the jaws are

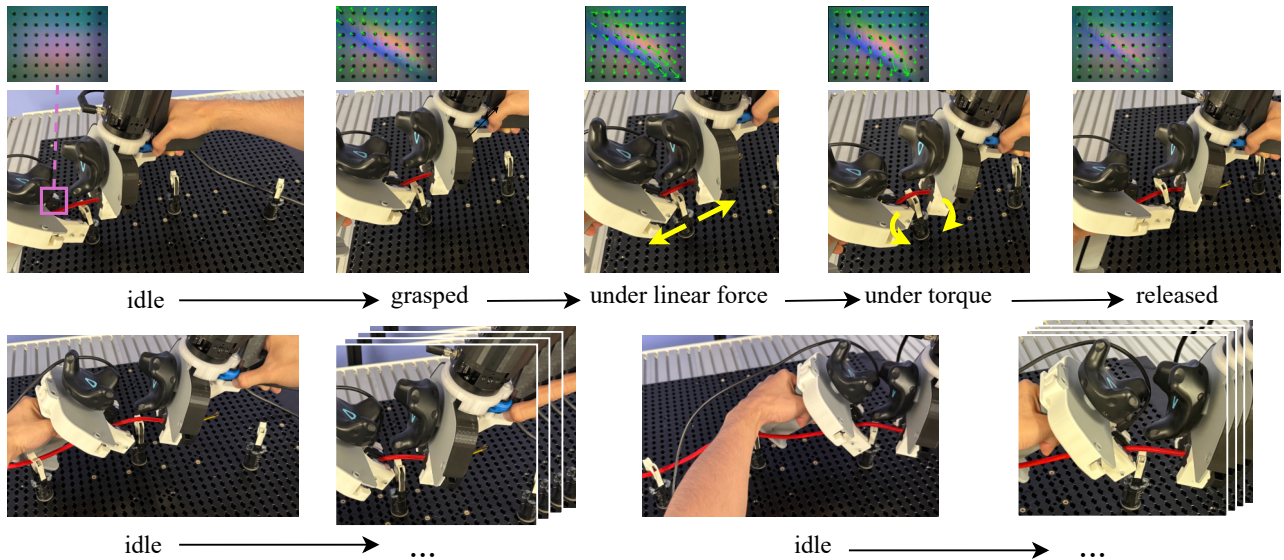


Fig. 5. **Overview of the cable mounting task process.** Using the proposed TacUMI, the cable is sequentially inserted into three U-type clips with different orientations.




Fig. 6. **Comparison of four gripper designs.** (a) Design without a locking mechanism: the fingertips are closed by pulling the trigger backward. (b) Design with a ratchet-based locking mechanism: once locked, the trigger can be released by pushing forward the two levers showed in the green circle. (c) Design with tension spring-driven fingertips: the fingertips are normally closed by a tension spring (showed in the green circle) and open when the trigger is pulled backward. (d) Design with a continuous locking mechanism: the proposed design introduced in detail in previous figure, enabling stable and adjustable locking at arbitrary grasping width.

closed only by continuously pulling the trigger. The second design (Fig. 6(b)) adds a ratchet-based locking mechanism, similar in principle to ForceMimic [16], allowing the trigger to be locked and released via mechanical levers. The third design (Fig. 6(c)) employs a reverse actuation structure with a tension spring mounted under the fingertips, keeping them normally closed, pulling the trigger opens the fingertips against the spring force. The fourth design (Fig. 6(d)) corresponds to our proposed continuous locking mechanism described in Section III-A, which ensures stable grasping during demonstrations.

We used each design to collect multimodal data during

the cable mounting task and mapped the measured F/T signals into the robot end-effector frame using Equations (3) and (4). As shown in Fig. 7(a), the first design success-




Fig. 7. **Comparison of force measurements in the cable mounting task across different gripper designs and teleoperated robot.** (a) Gripper without locking mechanism. (b) Gripper with ratchet-based locking mechanism. (c) Gripper with tension spring-based mechanism. (d) Proposed TacUMI gripper integrating tactile and force sensing. (e) Teleoperated robot with wrist-mounted FT sensor.

fully detected F/T data, however, since the operator had to continuously hold the trigger, an additional non-constant actuation force was introduced, which could not be removed through preprocessing. In Fig. 7(b), the second design with a ratchet-based locking mechanism exhibited discontinuous locking of the grasping width. As a result, the gripper failed to securely hold the cable, causing it to slip out during tightening and leading to unsuccessful task execution. In contrast, both Fig. 7(c) and Fig. 7(d) show F/T data that are consistent with those collected from the teleoperation in Fig. 7(e). Nevertheless, the spring-based design in Fig. 6(c)

lacks generalizability: due to the stiffness of the tension spring, it cannot grasp cables with larger diameters or objects requiring wider fingertip openings. In contrast, our final design TacUMI shown in Fig. 6(d) not only enables grasping across a wide range of grasping widths but also produces F/T data, that (after preprocessing) closely match those obtained from teleoperated robot with a wrist-mounted F/T sensor. The minor deviations are attributable to execution differences between handheld and teleoperated demonstrations. This confirms that the TacUMI design achieves correct and robot-consumable F/T data collection.

Beyond correctness, we also compare the efficiency of teleoperation with our handheld gripper design. In each trial of the cable mounting task, the cable had to be sequentially inserted into three clips arranged in different orientations. We measured the total task completion time across multiple trials for both methods. Results show that the teleoperation approach required an average of **4 minutes** per trial, whereas our handheld gripper completed the same task in only **1 minute 10 seconds** on average. This demonstrates a substantial improvement in data collection speed, enabling faster acquisition of high-quality multimodal datasets for long-horizon task segmentation.

C. Ablation on Model Architectures and Input Modalities

1) *Evaluation Metrics:* We evaluate our event segmentation models using the following metrics:

- **Frame-wise Accuracy:** The proportion of correctly predicted labels over all frames in the restored sequence after applying the soft voting strategy.
- **F1 Score per Class:** Harmonic mean of precision and recall for each skill class.

We conduct an extensive ablation study to analyze the impact of model architectures and input modalities on segmentation performance. For architectures, we compare BiLSTM, TCN, and Transformer encoder introduced in Sec. III-B. For input modalities, we evaluate single-modality settings, dual-modality combinations, three-modality combinations, and the full multimodal configuration (Table I).

TABLE I
FRAME-WISE ACCURACY (%) ON TACUMI DATA.

Input Modality	BiLSTM	TCN	Transformer
Camera only	0.7608	0.7217	0.3180
Camera + Tactile	0.9076	0.8880	0.6765
Camera + F/T	0.8632	0.8325	0.6645
Camera + Pose	0.8165	0.7675	0.4242
Camera + Tactile + F/T	0.9359	0.9051	0.7459
Camera + Tactile + F/T + Pose	0.9402	0.8945	0.7596

We first analyze the frame-wise accuracy results. BiLSTM consistently outperforms TCN and Transformer, as it can exploit both past and future context to accurately detect event boundaries and exhibits greater robustness to limited data scale (30k frames) and sensor noise, whereas the Transformer encoder requires larger datasets and is more sensitive to noise, and the TCN lacks the ability to access bidirectional temporal context. In terms of input modalities, models relying solely on third-person visual input perform worst, often

failing to identify subtle skill transitions. Adding tactile or F/T signals substantially boosts accuracy, while TCP pose alone provides only marginal gains. The best performance is obtained when combining all four modalities, with BiLSTM achieving the highest overall accuracy.

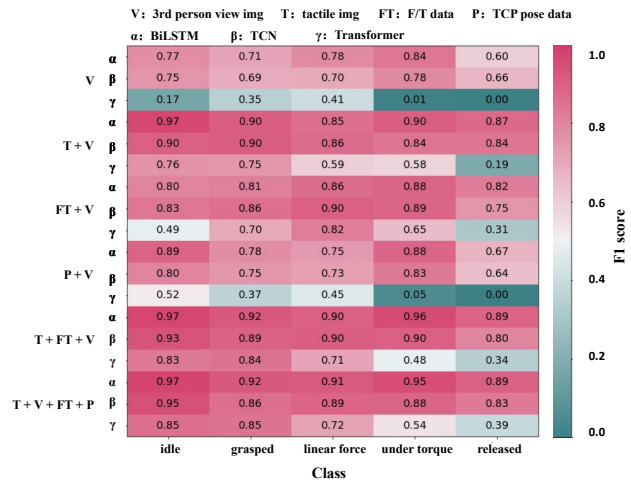


Fig. 8. **Per-class segmentation performance across input modalities and model architectures.** Heatmap of average F1-scores for each action class (*idle*, *grasped*, *under a linear force*, *under torque*, *released*) across all modality-architecture combinations. Higher values (towards red) indicate better class-wise segmentation accuracy.

Finally, we analyze class-wise segmentation performance (Fig. 8). For *grasped* and *idle* phases, most models achieve high F1 scores, as these skills involve clear and distinguishable signals. However, performance drops notably for the *released* and *under torque* classes, as the transitions into these states are subtle and less distinguishable from neighboring phases. The F1 score trends are consistent with the frame-wise accuracy results: performance improves as more modalities are incorporated, with BiLSTM consistently outperforming TCN and Transformer. Nevertheless, the improvement from three modalities (*third-person view* + *tactile* + *F/T*) to the full multimodal configuration (adding TCP pose) is negligible. This outcome is expected for our cable mounting task, since TCP pose remains nearly unchanged during critical phases such as grasping, stretching the cable (*under linear force*), and inserting it into the clip (*under torque*), and thus provides little additional information for segmentation. Among all events, the *released* class shows the lowest F1 score. This is primarily due to its extremely short duration, typically only 2–5 frames, so misclassifying even one or two frames leads to a sharp drop in F1 score.

In particular, BiLSTM with the full multimodal input yields balanced performance across all classes, demonstrating that complementary sensing modalities are critical for capturing fine-grained manipulation events.

D. Cross-Platform Validation on Robot Data

To verify that multimodal demonstrations collected with the handheld gripper can be reliably transferred to robot platforms, we additionally evaluate the trained segmentation models on teleoperation-collected datasets. Specifically, all

models are trained exclusively on **TacUMI** collected data but tested on robot-side demonstrations recorded via teleoperation.

As shown in Table II, when using only third-person camera input, the segmentation performance is very poor. This is expected, since the test set consists of robot end-effector manipulations, whereas the training and validation sets contain demonstrations with the handheld TacUMI, leading to a substantial domain gap between the two embodiments. Similar to the trends observed in Table I, adding tactile or F/T signals significantly improves accuracy; however, the results remain 10–20% lower than those achieved with gripper-collected test sets. This gap can be explained by the higher difficulty of controlling two robots via *sigma.7* haptic devices, which introduces additional noise and less stable demonstrations compared to handheld execution. Importantly, when tactile and FT modalities are combined, the accuracy nearly matches that of the gripper test set, further confirming the critical role of multimodal fusion in achieving robust event segmentation across platforms.

TABLE II
FRAME-WISE ACCURACY (%) ON ROBOT DATA

Input Modality	<i>BiLSTM</i>	<i>TCN</i>	<i>Transformer</i>
Camera only	0.2288	0.1820	0.2227
Camera + Tactile	0.7474	0.6002	0.5351
Camera + F/T	0.6611	0.6366	0.4126
Camera + Pose	0.4694	0.4636	0.2256
Camera + Tactile + F/T	0.9104	0.8793	0.8092
Camera + Tactile + F/T + Pose	0.9155	0.7262	0.7796

In conclusion, these results confirm that data collected with TacUMI not only supports accurate event segmentation in handheld demonstrations, but also transfers effectively to robot-executed tasks. This cross-platform generalization highlights the correctness of our data collection design and its suitability for downstream robot learning.

V. CONCLUSION

We presented TacUMI, a handheld interface that extends the Universal Manipulation Interface with multimodal sensing and task segmentation. By integrating fingertip tactile sensing, a wrist-mounted force–torque sensor, and drift-free 6-DoF pose tracking, it enables high-quality demonstrations for contact-rich, long-horizon tasks. A continuous locking mechanism ensures reliable force measurements and stable grasps. We further introduced a sequence modeling framework that segments multimodal demonstrations into reusable skills. Experiments on cable mounting demonstrate accurate segmentation and reliable transfer to robot execution.

Although these results validate the hardware and segmentation approach, broader evaluation is needed to assess generalizability across tasks. Beyond segmentation, TacUMI can support multimodal imitation learning, enabling robots to acquire executable skills directly from demonstrations and advancing toward end-to-end learning for complex manipulation.

ACKNOWLEDGMENT

The authors acknowledge the financial support by the Bavarian State Ministry for Economic Affairs, Regional De-

velopment and Energy (StMWi) for the Lighthouse Initiative KI.FABRIK (Phase 1: Infrastructure as well as the research and development program under grant no. DIK0249). The authors acknowledge the financial support by the Fundamental Research Funds for the Central Universities (grant no. KG202514).

REFERENCES

- [1] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” 2023.
- [2] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, p. 02783649241273668, 2023.
- [3] R. Wolf, Y. Shi, S. Liu, and R. Rayyes, “Diffusion models for robotic manipulation: A survey,” *arXiv preprint arXiv:2504.08438*, 2025.
- [4] Y. Zhu, P. Stone, and Y. Zhu, “Bottom-up skill discovery from unsegmented demonstrations for long-horizon robot manipulation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4126–4133, 2022.
- [5] M. Guo and M. Bürger, “Geometric task networks: Learning efficient and explainable skill coordination for object manipulation,” *IEEE Transactions on Robotics*, vol. 38, no. 3, pp. 1723–1734, 2021.
- [6] K. Chen, Z. Shen, Y. Zhang, L. Chen, F. Wu, Z. Bing, S. Haddadin, and A. Knoll, “Lemmo-plan: Llm-enhanced learning from multimodal demonstration for planning sequential contact-rich manipulation tasks,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 11972–11978.
- [7] K. Chen, Z. Bing, F. Wu, Y. Meng, A. Kraft, S. Haddadin, and A. Knoll, “Contact-aware shaping and maintenance of deformable linear objects with fixtures,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 1–8.
- [8] K. Shaw, S. Bahl, and D. Pathak, “Videodex: Learning dexterity from internet videos,” in *Conference on Robot Learning*. PMLR, 2023, pp. 654–665.
- [9] Y. Zhu, A. Joshi, P. Stone, and Y. Zhu, “Viola: Imitation learning for vision-based manipulation with object proposal priors,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1199–1210.
- [10] J. Campbell and K. Yamane, “Learning whole-body human-robot haptic interaction in social contexts,” in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 10177–10183.
- [11] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song, “Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [12] Zhaxizhuoma, K. Liu, C. Guan, Z. Jia, Z. Wu, X. Liu, T. Wang, S. Liang, P. Chen, P. Zhang, H. Song, D. Qu, D. Wang, Z. Wang, N. Cao, Y. Ding, B. Zhao, and X. Li, “Fastumi: A scalable and hardware-independent universal manipulation interface with dataset,” 2025.
- [13] M. Xu, Z. Xu, C. Chi, M. Veloso, and S. Song, “Xskill: Cross embodiment skill discovery,” in *Conference on robot learning*. PMLR, 2023, pp. 3536–3555.
- [14] D. Tanneberg, K. Ploeger, E. Rueckert, and J. Peters, “Skid raw: Skill discovery from raw trajectories,” *IEEE robotics and automation letters*, vol. 6, no. 3, pp. 4696–4703, 2021.
- [15] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, “Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 12156–12163.
- [16] W. Liu, J. Wang, Y. Wang, W. Wang, and C. Lu, “Forcemimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 1105–1112.
- [17] T. Eiband, J. Liebl, C. Willibald, and D. Lee, “Online task segmentation by merging symbolic and data-driven skill recognition during kinesthetic teaching,” *Robotics and Autonomous Systems*, vol. 162, p. 104367, 2023.
- [18] S. B. Kang and K. Ikeuchi, “Toward automatic robot instruction from perception-temporal segmentation of tasks from human hand motion,” *IEEE Transactions on Robotics and Automation*, vol. 11, no. 5, pp. 670–681, 1995.