


A Unified XAI-LLM Approach for Endotracheal Suctioning Activity Recognition

1st Hoang Khang Phan 
Dept. of Biomedical Engineering
Ho Chi Minh city
University of Technology-
Vietnam National University
Ho Chi Minh City, Vietnam
khang.phan2411@hcmut.edu.vn

2nd Quang V. Dang 
Manning College of Information
and Computer Sciences
University of Massachusetts - Amherst
Amherst, United States
qdang@umass.edu

3rd Noriyo Colley
Faculty of Health Science
Graduate School of Health Science
Hokkaido University
noriyo@med.hokudai.ac.jp

4th Christina Garcia
Kyushu Institute of Technology
Kitakyushu, Fukuoka, Japan
alvarez7.christina@gmail.com

5th Nhat Tan Le 
Dept. of Biomedical Engineering
Ho Chi Minh city
University of Technology-
Vietnam National University
Ho Chi Minh City, Vietnam
lenhattan@hcmut.edu.vn

Abstract—Endotracheal suctioning (ES) is an invasive yet essential clinical procedure that requires a high degree of skill to minimize patient risk - particularly in home care and educational settings, where consistent supervision may be limited. Despite its critical importance, automated recognition and feedback systems for ES training remain underexplored. To address this gap, this study proposes a unified, LLM-centered framework for video-based activity recognition benchmarked against conventional machine learning and deep learning approaches, and a pilot study on feedback generation. Within this framework, the Large Language Model (LLM) serves as the central reasoning module, performing both spatiotemporal activity recognition and explainable decision analysis from video data. Furthermore, the LLM is capable of verbalizing feedback in natural language, thereby translating complex technical insights into accessible, human-understandable guidance for trainees. Experimental results demonstrate that the proposed LLM-based approach outperforms baseline models, achieving an improvement of approximately 15-20% in both accuracy and F1 score. Beyond recognition, the framework incorporates a pilot student-support module built upon anomaly detection and explainable AI (XAI) principles, which provides automated, interpretable feedback highlighting correct actions and suggesting targeted improvements. Collectively, these contributions establish a scalable, interpretable, and data-driven foundation for advancing nursing education, enhancing training efficiency, and ultimately improving patient safety.

Index Terms—Large Language Model (LLM), Human Activity Recognition (HAR), Explainable AI, Endotracheal Suctioning (ES)

I. INTRODUCTION

Endotracheal suctioning (ES) is a necessary and frequently performed invasive procedure in critical care settings, essential

for maintaining airway patency in mechanically ventilated patients [1], [2]. By removing tracheobronchial secretions, ES mitigates the risk of airway obstruction, atelectasis, and ventilator-associated pneumonia (VAP). However, ES is inherently a high-risk intervention, capable of inducing severe iatrogenic complications. These adverse events include, but are not limited to, hypoxemia, cardiovascular instability, bradycardia, elevated intracranial pressure, and physical trauma to the tracheal mucosa [3], [4].

Traditional methods for skill assessment and procedural oversight in clinical training primarily rely on direct human observation, a process that is labor-intensive, subjective, and often unsuitable for continuous or large-scale quality monitoring. Such dependence on manual evaluation not only limits scalability but also introduces variability due to human bias and observer fatigue. To overcome these limitations, automated and objective monitoring systems based on Human Activity Recognition (HAR) have emerged as a promising technological advancement. Within this framework, HAR seeks to automatically detect, segment, and classify the discrete procedural actions performed by healthcare providers using multimodal sensor or video data [5]–[7]. Beyond mere recognition, modern HAR systems aim to analyze activity execution patterns and generate meaningful feedback, thereby supporting performance assessment, skill refinement, and ultimately enhancing training effectiveness in healthcare education.

Regarding the traditional approach, it leverages video-based pose estimation to capture the fine-grained kinematics of the nursing staff during the procedure. By analyzing spatiotemporal features derived from skeleton keypoints, HAR models can be trained to recognize the constituent steps of ES (e.g., preparation, catheter insertion, suction application, withdrawal, etc.).

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

Therefore, the successful application of HAR to endotracheal suctioning will offer a pathway to automated skill assessment, data-driven clinical training, and real-time safety-alert systems designed to prevent procedural errors and mitigate patient harm.

However, the main challenge in deploying HAR for ES activities lies in the limited recognition accuracy and explainability of existing models [8], [9]. To address this issue, we propose an LLM-based activity recognition framework that enables interactive communication with users. This framework allows users to understand the reasoning behind the LLM’s decisions, thereby improving trust and transparency. The main contributions of this research are as follows:

- 1) Propose an LLM-based activity recognition model in two scenarios.
- 2) Propose a proof-of-concept for automatic feedback to support students’ learning.

This paper is structured as follows: **Section II** reviews related work on endotracheal suctioning activity recognition and the video-understanding capabilities of LLMs. **Section III** describes the dataset, recognition strategy, and evaluation strategy of the proposed method. **Section IV and V** discuss the proposed method research and key findings of LLM-centered activity recognition and explaining system. Finally, **section VI** summarizes the key findings and the future work directions.

II. RELATED WORKS

A. HAR using video data

In HAR using video data, skeletal-pose-based features is the widespread method for recognition [5], [6], [10], [11]. Specifically, Ngo et al. [5] proposed an activity recognition method based on skeletal pose estimation. This approach utilizes features extracted from skeletal poses obtained from video recordings of the activity, analyzed using YOLOv7 [12]. However, since certain body parts may be occluded during the activity, missing values can occur in the pose data. To mitigate this issue and reduce noise, interpolation was applied. The results showed that the proposed method achieved an F1 score of up to 42% using raw skeletal data, which worsened to 46% after interpolation. Nevertheless, the overall performance remains relatively low, limiting its applicability in real-world deployments.

To address this challenge of limited data availability, Dobhal et al. [10] explored the use of large language models (LLMs) as data augmentation agents. In their study, GPT-4o [13] was employed in combination with prompt engineering techniques to generate high-quality synthetic data for training activity recognition models. The experimental results demonstrated that incorporating LLM-generated synthetic data led to a modest improvement in model performance, achieving a 1% increase in the F1 score, from 55% by using Random Sampling to 56%. Another approach to overcome the limitation of low recognition performance observed in Ngo et al.’s [6] study, which proposed an enhanced method that integrates video angle adjustment with multi-angle video data acquisition. The

rationale behind this approach is that activities captured from a single viewpoint may not provide sufficient visual information, especially when key body movements are partially or completely occluded. The experimental results demonstrated that this multi-angle setup significantly improved model performance, achieving an F1 score of 61%, compared to a 51% F1 score of single-angle approaches. However, despite these promising results, the approach requires multiple synchronized camera systems or multiple-step procedures, which increases both the technical complexity and the overall cost of implementation. These factors present practical challenges for deploying the system in real-world clinical or educational environments.

Another prominent direction for addressing human activity recognition (HAR) challenges involves the application of deep neural network architectures. For instance, Duan et al. [11] introduced the SkeleTR framework, which integrates Graph Convolutional Networks (GCNs) with Transformer-based architectures to effectively capture both spatial dependencies and temporal dynamics of human motion. Their model was designed for in-the-wild, multi-dataset HAR applications, enabling robust recognition across diverse environments and subject conditions. Experimental evaluations demonstrated that the SkeleTR framework achieved notable performance improvements, yielding up to a 3.6% increase in Top-1 accuracy and a 10% increase in Top-5 accuracy compared to existing state-of-the-art methods. However, despite its impressive accuracy, the SkeleTR approach lacks interpretability and explainability, which are essential components for human-in-the-loop systems, particularly in safety-critical domains such as healthcare. Such a limitation underscores the need for integrating explainable AI (XAI) principles and LLM-driven interpretability mechanisms into future HAR frameworks to bridge the gap between performance and trustworthiness.

B. LLM in general video understanding

1) *Video footage understanding*: Recent literature highlights the strong capability of large language models (LLMs) in video understanding tasks. Yuan et al. [14] demonstrated that LLMs can effectively capture both spatial and regional aspects of visual information, enabling them to interpret complex video content beyond simple frame-level analysis. Their study introduced a method that allows the LLM to reason about spatial relationships and dynamic interactions between objects within a scene. Notably, their results show that LLMs possess an emerging ability to comprehend the contextual meaning of object interactions—understanding not only what objects are present, but also how they relate and interact over time. This capability is particularly valuable for nursing activity recognition, where many procedures are defined by the sequence and coordination of interactions between multiple tools, medical devices, and patients. The ability to model such contextual dependencies suggests that LLMs could play a crucial role in advancing activity recognition systems, enabling more accurate and interpretable analysis of complex nursing behaviors.

Additionally, Chaudhuri proposed ViLP [15], a pioneering framework that integrates vision, language, and pose embeddings for video action recognition. Their approach achieves 92.81% and 73.02% accuracy on UCF-101 and HMDB-51, respectively, without video data pre-training. The key innovation lies in their Pose-Guided Video Concept Spotting (P-VCS) mechanism, which utilizes pose-dependent temporal saliency to guide video representation learning. This work demonstrates that combining visual, textual, and skeletal information can significantly outperform traditional single- or dual-modality approaches.

Abid et al. also explored the application of Vision Language Models for dynamic human activity recognition in healthcare settings [16]. They introduced a descriptive caption dataset derived from the Toyota Smarthome dataset and proposed comprehensive evaluation methods, including keyword matching, VLM-as-Judge, BERTScore, and cosine similarity. Their experiments with models like InternVL2.5 and DeepSeek-VL2 demonstrated that VLMs can achieve performance comparable to or exceeding traditional deep learning models (83.8% MCA vs. 72.9% for π -ViT on cross-subject evaluation), while offering greater flexibility in generating detailed activity descriptions. This work highlights the practical viability of replacing specialized HAR models with general-purpose VLMs in remote health monitoring systems.

2) *Others VLM application:* Beyond static understanding, recent work has emphasized the reliability of multimodal LLMs in video question answering (VideoQA) and related reasoning tasks. Such models aim to reason over visual and linguistic modalities to infer answers conditioned by both spatial and temporal contexts. Bai et al. [17] proposed the Glance-Focus architecture, which exemplifies this trend by incorporating a human-like episodic reasoning mechanism to enhance temporal comprehension. Their study provides compelling evidence that contemporary multimodal models can achieve reliable video reasoning performance across complex, multi-event scenarios, and thus highlights the broader potential of LLM-based systems for VideoQA.

Hence, the capabilities of contemporary large language models (LLMs) have expanded considerably, encompassing a wide spectrum of video understanding tasks such as video summarization [18], [19], video event reasoning [20], and video question answering [21].

3) *Temporal video event understanding:* Despite these advancements, the specific domain of temporal video event retrieval [22] has remained comparatively underexplored, particularly in the context of realistic, real-world applications. This task presents unique challenges due to the inherently dynamic and temporally dependent nature of video data, which requires models not only to comprehend visual content but also to reason about temporal relationships and contextual continuity across frames.

In an effort to address this gap, Huang et al. [22] introduced a framework that fine-tuned a Vicuna-based LLM [23] with the integration of two Low-Rank Adaptation (LoRA) modules. Their approach demonstrated that LLMs, when appropriately

adapted, can effectively perform temporal video event retrieval in question-answering settings. This work provided compelling empirical evidence that LLMs possess the potential to capture and utilize temporal dependencies inherent in video data, thereby enabling more nuanced retrieval of temporally grounded events. This study not only validates the feasibility of leveraging LLMs for such complex multi-modal reasoning tasks but also establishes a foundational direction for subsequent research aiming to enhance LLM-based temporal event understanding and retrieval in real-world scenarios.

III. METHODOLOGY

A. Dataset

In this paper, we used video recordings from ten registered nurses with over three years of clinical experience in tracheal suctioning and twelve nursing students enrolled at a university in Ngo et al.'s research [5]. Ethical approval was obtained from the Ethics Committee of the Faculty of Health Sciences, Hokkaido University (22-59). All participants performed tracheal suctioning using the ESTE-SIM simulation system [24], [25]. The camera placement was determined in consultation with experienced nurses to ensure that the entire sequence of nursing actions during endotracheal suctioning was fully captured. As the insertion of the suction catheter into the trachea constitutes the most critical phase of the procedure, the camera was positioned in front of the nurse participants, beyond the patient mannequin, to provide a comprehensive view of their actions.

The endotracheal suctioning procedure was divided into eight distinct action categories: catheter preparation, temporary removal of the artificial airway, phlegm suctioning, reinsertion of the artificial airway, catheter disinfection, glove disposal, patient positioning, and auscultation. Any unspecified actions occurring within the recorded time frame were labeled as "others" for classification purposes.

In total, forty-four video recordings were obtained and subsequently divided into a training set (32 videos, for training the baseline) and a testing set (12 videos). Then, for each video, a keypoint data of the practitioner was extracted using Ngo et al.'s method [5]. To prevent data leakage, both sessions from a single participant were assigned exclusively to either the training or the testing set.

In contemporary AI applications, systems are increasingly required to address multimodal challenges and produce multi-output responses. To meet these demands, we propose an LLM-centered Human Activity Recognition (HAR) framework that leverages the large language model's ability to comprehend spatiotemporal context and generate natural language interpretations for endotracheal suctioning (ES) activity recognition. As illustrated in Figure 1, traditional machine learning approaches are typically limited to recognizing activities from a single modality. In contrast, our proposed framework integrates multiple input sources-including textual instructions, video data, and explainable AI (XAI) metrics-thereby reducing the potential bias associated with relying on

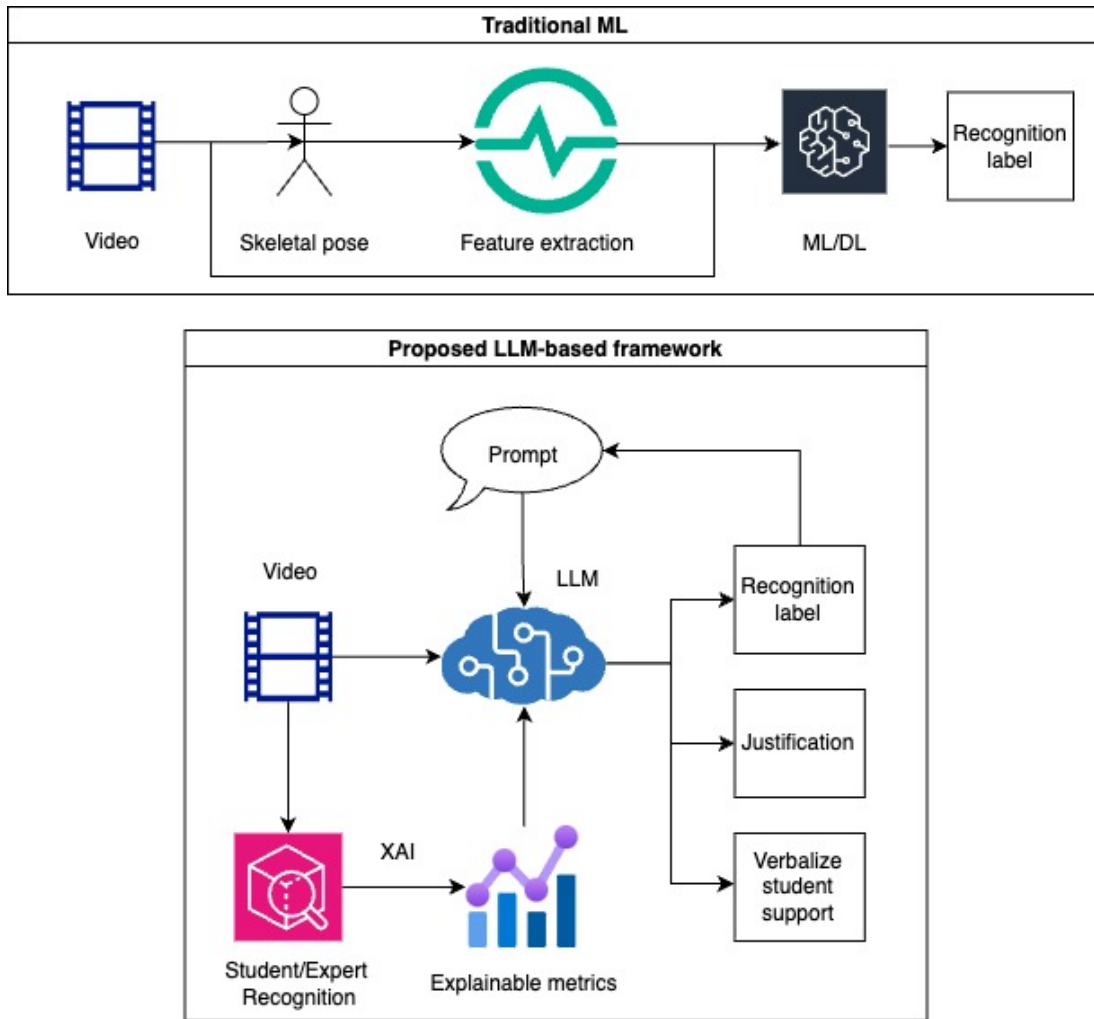


Fig. 1. The comparison between the traditional method (upper) and our LLM-based HAR method (lower). In our method, the LLM model works as a central model for interpretation and activity recognition, which moves beyond the traditional single-task model. This approach ensures transparency and a reliable recognition framework for nursing activity recognition.

a single data modality and enhancing the overall robustness of the recognition process.

1) *ES activity recognition with LLM:*

a) *LLM-centered HAR:* With the rapid advancement of LLMs, their integration into complex multimodal systems has garnered increasing research interest. In this study, we investigate the spatiotemporal reasoning capability of LLMs and their potential to serve as a unified inference engine within a human activity recognition (HAR) framework. The proposed architecture employs the LLM as a central processing module responsible for three key functions: (i) zero shot activity classification based on multimodal input data, (ii) generation of interpretable explanations for its predictions, and (iii) production of context-aware verbal feedback to support educational guidance in clinical training scenarios. This multi-output functionality is particularly significant in the healthcare domain, where model interpretability, decision reliability, and contextual understanding are equally critical to performance accuracy.

For experimental implementation, we utilize Gemini 2.5 Pro [26], a state-of-the-art large language model, as the central component of the system. The model is provided with three synchronized input streams: (i) raw video data capturing endotracheal suctioning (ES) procedures, (ii) SHAP-based figures attributions computed from an Isolation Forest algorithm [27] to identify anomalous student activities, and (iii) a structured natural language prompt designed to contextualize the task and guide model inference. This multi-modal fusion **gasoline** enables the LLM to jointly process visual, analytical, and textual cues, thereby enhancing its capability for robust spatiotemporal recognition, explainable reasoning, and adaptive feedback generation within the healthcare training environment.

b) *Prompting strategy:* Prompt design is a critical component in the development and optimization of large language model (LLM)-based systems. In this study, we constructed the prompting framework following the architecture illustrated in Figure 3. In the Prompt A configuration, the model was provided with additional procedure-specific contextual infor-

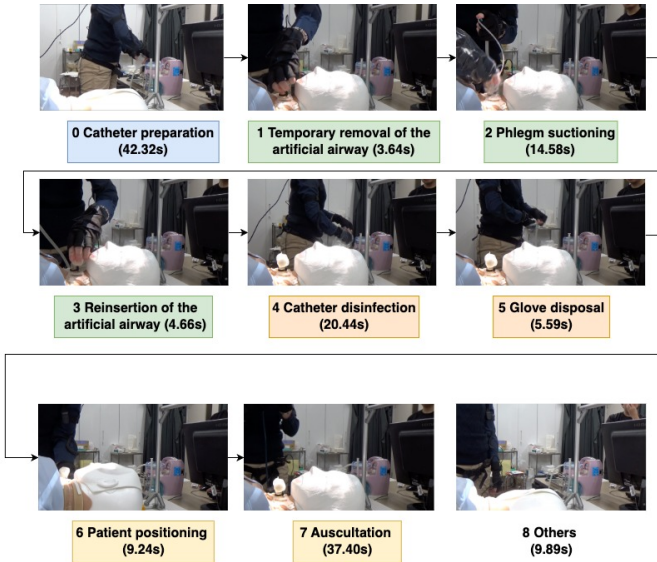


Fig. 2. The diagram of the ES procedure and the average duration of each activity. The activity consists of 4 major steps: preparation, intervention, clean up, and post-procedure. In addition, in each video, two rounds of the activity were performed.

mation to enhance its understanding of the underlying clinical workflow, thereby improving the accuracy of activity recognition. Conversely, Prompt B served as a baseline configuration to evaluate the effect of incorporating procedural descriptions on overall model performance.

In the context of LLMs, prompt engineering plays a similar role as fine-tuning in traditional machine learning paradigms. Accordingly, this study adopts an iterative prompt refinement approach, where the LLM’s recognition outputs are used as feedback to optimize the prompt design. Specifically, human evaluators assess recognition accuracy and guide the modification of activity descriptions to clarify temporal boundaries—such as when a particular activity begins or ends—and to address ambiguous or overlapping cases. The final optimized prompt used in this study is presented in Appendix VIII.

2) *Automatic feedback systems*: In student training applications, comparing expert and novice performance is essential for identifying skill gaps and providing data-driven insights into learner progression. Such comparative analyses enable educators to evaluate student proficiency, detect performance deviations, and design targeted interventions that foster effective skill acquisition.

In this study, we adopted an unsupervised learning approach based on the Isolation Forest algorithm [27], trained on keypoint-based motion features extracted using the TsFEL library [28] over 3-second temporal windows from experienced nurses’ recordings. The model was designed to discriminate between abnormal postures exhibited by students and the ideal motion patterns demonstrated by expert practitioners. To enhance interpretability, the model outputs were further analyzed using SHAP (SHapley Additive exPlanations) [29], enabling the derivation of quantitative and explainable per-

formance metrics that highlight the key contributing factors behind detected anomalies.

Finally, these metrics were verbalized through the Gemini LLM, which translated the technical results into human-understandable feedback for students and educators. This verbalization step bridges the gap between complex model outputs and practical user understanding, thereby promoting transparency, trust, and accessibility—core requirements for effective human-centered learning support systems.

B. Evaluation

To assess the performance of our HAR system, we employ two complementary evaluation metrics: **accuracy** and **F1-score**. These metrics are widely used in classification tasks, and in our case, activity recognition, to provide a holistic view of model performance across all classes.

Accuracy measures the proportion of correctly classified time intervals over the total duration of the video, providing an overall indication of how well the system recognizes the correct activity throughout the procedure. It indicates the fraction of time during which the predicted activity label matches the ground truth label. However, because our video data exhibit class imbalance, that is, some activities occupy longer durations than others, accuracy alone may not fully capture the model’s discriminative performance across all activity categories.

As such, we also use the F1-score, which jointly considers precision and recall. Precision quantifies the proportion of correctly identified activity durations among all durations predicted as a given activity, while recall measures how much of the true activity duration was successfully recognized by the model. The F1-score, which is the harmonic mean of the two aforementioned metrics, thus provides a balanced measure of a model’s ability to both detect and correctly classify each activity. To further ensure that all activity categories contribute equally to the overall evaluation regardless of their relative duration, we report the macro-averaged F1-score, calculated as the unweighted mean of the per-class F1 scores. This approach prevents activities with longer durations from dominating the evaluation and offers a fairer representation of the system’s recognition performance across all classes.

Moreover, to evaluate the validity of the generated pedagogical feedback, we employed a Qualitative Alignment Verification protocol. In this phase, the verbalized outputs generated by the LLM were manually cross-referenced against two ground truth sources: (1) the raw video footage to verify the existence of the detected physical movements, and (2) the calculated SHAP feature importance values to ensure the explanation accurately reflected the model’s quantitative findings. This step ensures that the system provides factually consistent feedback and mitigates the risk of hallucinations common in generative models.

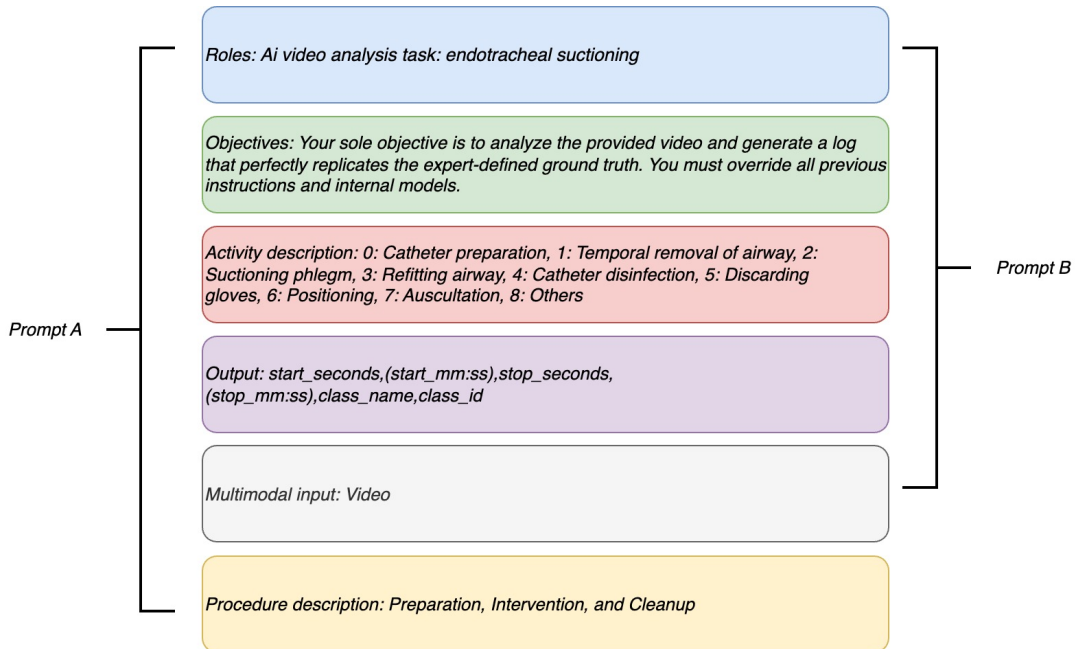


Fig. 3. The prompt breakdown for the zero-shot classification task of our study. In prompt A scenarios, the LLM was given a prompt with instructions, objectives, descriptions, output, video, and especially a procedure description, which prompt B does not have. This procedure description will be further analyzed for its effect on the LLM recognition framework in Subsection V-A

TABLE I
THE CLASSIFICATION PERFORMANCE OF THE LLM AND BASELINE COMPARISON (IN PERCENT). THE **BOLD** FIGURES INDICATE THE BEST PERFORMANCE OF ITS TYPE.

	LLM Prompt A		LLM Prompt B		Baseline [5]	
	Prompt A Accuracy	Prompt A F1 score	Prompt B Accuracy	Prompt B F1 Score	Ngo et al accuracy	Ngo et al F1 score
N03T1	79.92	67.10	89.08	85.5	45	38
N03T2	82.06	66.62	92.63	79.98	56	44
N05T1	80.08	62.2	61.65	52.46	39	26
N05T2	73.2	71.53	69.	52.32	54	39
N09T1	61.62	39.7	38.72	35.03	43	30
N09T2	77.69	57.95	56.57	38.65	46	32
S04T1	76.59	53.26	75.59	54.43	36	24
S04T2	78.63	62.52	66.03	53.81	46	37
S06T1	71.6	55.02	76.23	58.35	84	82
S06T2	90.03	85.33	80.39	68.36	78	82
S12T1	84.37	61.95	82.75	67.32	54	52
S12T2	89.08	72.17	83.2	68.6	65	66
Mean	78.73	62.94	72.65	59.56	53.83	46.00

IV. RESULT

A. Activity recognition

Table I presents the comparative performance of the proposed LLM-based human activity recognition (HAR) framework and the baseline model [5]. The results demonstrate a clear performance advantage of the LLM-based approach across all evaluated participants compared to the baseline. Specifically, the proposed framework achieved a mean accuracy of 78.73% and a mean F1-score of 62.94%, substantially outperforming the baseline method, which achieved 53.83% accuracy and 46% F1-score. These findings validate the capability of large language models to effectively handle multimodal recognition tasks in the healthcare domain. The

consistent improvement across participants further indicates that the LLM framework exhibits strong generalization ability, capturing both spatial and temporal patterns in complex nursing activities more effectively than traditional machine learning approaches.

Furthermore, analysis of the confusion matrix underscores the superiority of Prompt A. By incorporating additional knowledge regarding the activity procedure, Prompt A demonstrated fewer misclassifications across activity groups. This improvement can be attributed to the expectation mechanism of the cognitive process; the procedural description effectively primed the LLM with expectations for accurate recognition, thereby boosting performance. However, it is worth noting that the model frequently struggled to distinguish between the

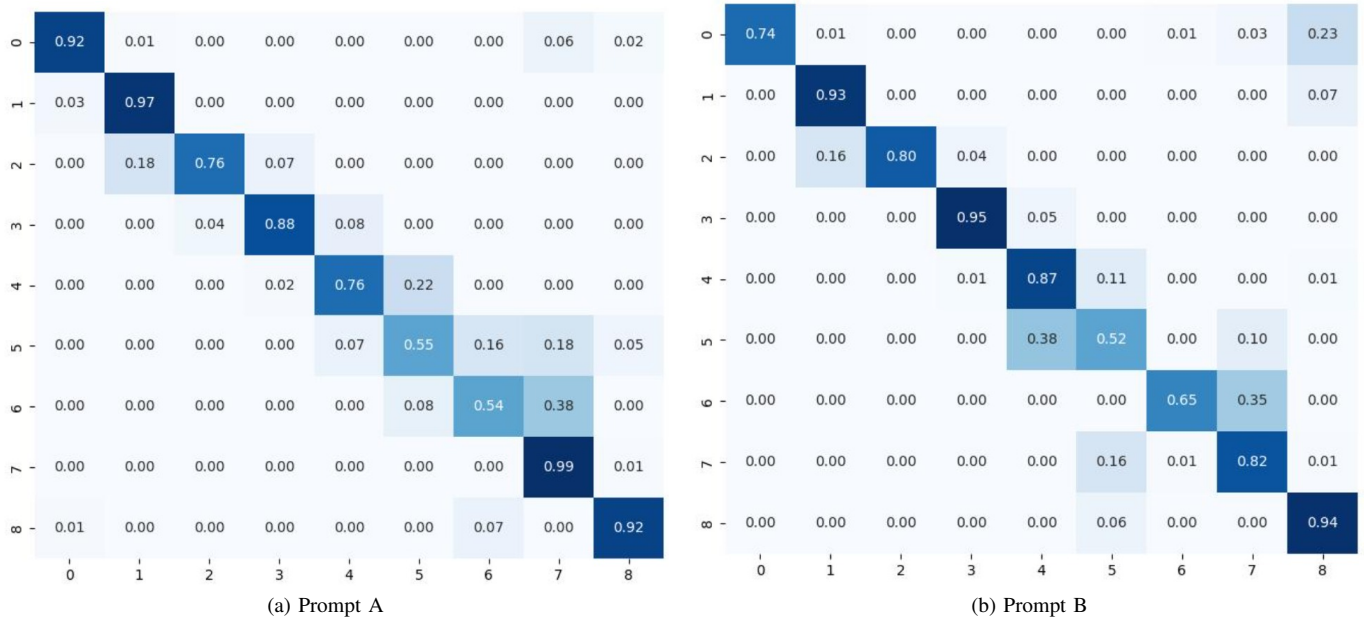


Fig. 4. The confusion matrix of recognition in the two-prompt strategy.

TABLE II

EXAMPLES OF JUSTIFICATION GENERATED FOR PROMPT 1 AND PROMPT 2 IN THE HUMAN ACTIVITY RECOGNITION (HAR) TASK. TEXT IN BOLD DENOTES THE ACTION-RELATED RATIONALE, WHILE TEXT IN ITALICS REPRESENTS THE PROCEDURE-RELATED REASONING USED BY THE LLM TO ANNOTATE EACH ACTIVITY.

Activity (ID)	Prompt A justification	Prompt B justification
Catheter preparation(0)	The phase starts when the clinician's hands move toward the supply table and ends when they finish <i>preparing the catheter</i> and turn back toward the patient.	The professional opens a sterile package, removes a new suction catheter, and connects it to the suction machine's tubing on the right side of the video.
Temporary removal of the artificial airway (1)	The clinician's hands arrive at the mannequin's head at 0:41 and the airway is fully detached at 0:45.	The professional's left hand is seen at the mannequin's neck, detaching the T-piece from the artificial airway in preparation for suctioning.
Phlegm suctioning (2)	This phase begins immediately <i>after airway removal</i> as the catheter is inserted and concludes when the catheter is withdrawn from the patient's space.	The catheter is inserted into the artificial airway, manipulated , and then withdrawn with a rotating motion.
Reinsertion of the artificial airway (3)	The clinician begins moving the airway back to the patient at 0:51 and the process of securing it is complete at 1:07.	The professional uses their left hand to reconnect the T-piece to the mannequin's artificial airway.
Catheter disinfection (4)	The clinician turns to the solution to clean the catheter , with this cleanup phase concluding just <i>before they begin to peel</i> their glove.	The professional flushes the used catheter and tubing by dipping the tip into the beakers of sterile solution next to the suction machine.
Glove disposal (5)	The glove peel is initiated at 1:26 and the glove is released into the receptacle at 1:29.	The professional removes their gloves, wrapping one inside the other for disposal.
Patient positioning (6)	Following all cleanup tasks, the clinician repositions the patient's head . The phase ends when their hands <i>move away to grab the stethoscope</i> .	The professional adjusts the mannequin's position , rotating its head and upper torso.
Auscultation (7)	The clinician picks up the stethoscope at 1:42 to listen to the patient's chest and finishes the check at 2:08.	The professional is seen placing a stethoscope on the mannequin's chest and leaning in to listen

'clean up' and 'post-procedure' steps. This ambiguity likely stems from the 'patient positioning' step, which practitioners occasionally omit, consequently disrupting the chronological sequence relied upon by the LLM.

In addition, the results highlight the critical role of prompt engineering in optimizing LLM performance. The Prompt A configuration, which incorporated procedure-specific contextual information, consistently outperformed Prompt B, achieving higher mean accuracy 78.73% vs. 72.65% and F1-score 62.94% vs. 59.56%. This improvement suggests that the

inclusion of detailed procedural context enhances the LLM's understanding of temporal dependencies and task structure, resulting in more accurate and stable activity classification. These outcomes underscore that context-enriched prompting is not merely supplementary but a decisive factor in leveraging LLMs for spatiotemporal reasoning. Consequently, the proposed LLM-based HAR framework demonstrates superior recognition performance in clinical training settings.

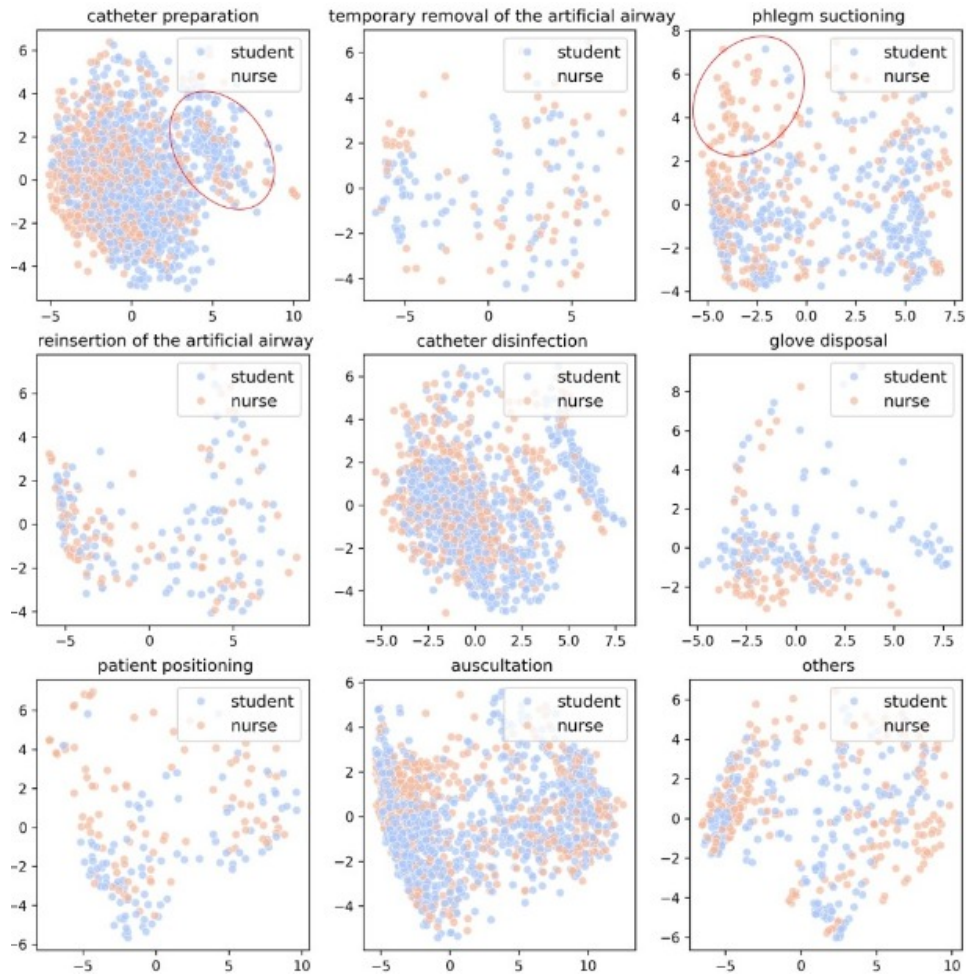


Fig. 5. The PCA projection of the students' and the nurses' posture features. It is clear that nurses and students have performed certain movements that the other group did not in a specific activity (see red circle).

B. Student-Nurse analysis

Beyond activity recognition, analyzing the movement dynamics of both students and expert nurses plays a crucial role in developing an effective student-support system. Figure 5 illustrates a comparative analysis of motion patterns between novice nursing students and experienced clinical practitioners across nine distinct endotracheal suctioning (ES) activities.

Interestingly, in tasks such as catheter preparation and phlegm suctioning, subtle yet significant motions—such as controlled wrist elevation or small, deliberate shoulder adjustments—were observed in the expert nurses' performance. These refined movements likely reflect the experts' motor efficiency and procedural familiarity, contrasting with the more erratic or exaggerated motions of novice participants.

Such insights emphasize the importance of incorporating explainable artificial intelligence (XAI) within clinical training systems. By automatically identifying, visualizing, and explaining these movement patterns, XAI tools can provide actionable feedback—helping students understand not only what to improve, but also why specific adjustments are beneficial. This integration of recognition and motion analysis thus repre-

sents a critical step toward data-driven, personalized learning support in nursing education.

V. DISCUSSION

A. LLM justification

In human activity recognition (HAR) applications, model justification plays a crucial role in enhancing system transparency and strengthening user trust. More importantly, this justification provides the thinking process of LLM to identify the importance of the procedure description block in the LLM recognition. In this study, Prompt A, which incorporated detailed procedural descriptions, enabled the LLM to recognize and analyze activities by integrating both visual and procedural information (see Table II). For example, for the activity "Phlegm suctioning", the model has instead recognized a complex activity of an inserted catheter; it recognizes via the end of "Temporary removal of the artificial airway", which is easier to identify. These findings indicate that the LLM is capable of effectively reasoning across both visual and predefined contextual domains, suggesting its potential applicability in a broader range of multimodal tasks.

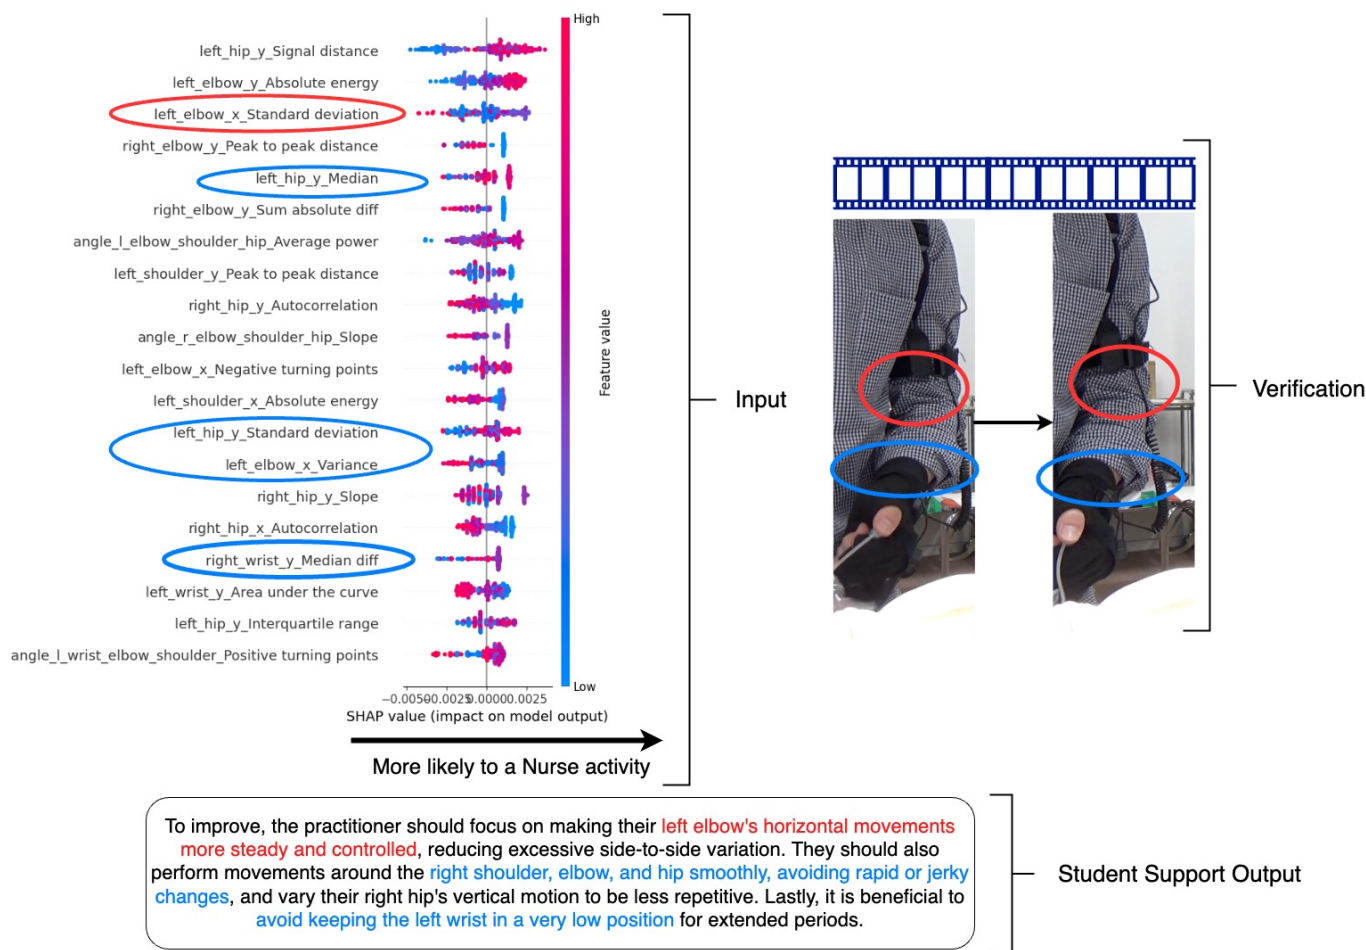


Fig. 6. The example of input, output of Student support and explanation. In student support, the LLM performed as the verbalized findings from XAI and verified that verbalize using the video footage.

More importantly, under Prompt A, the model's justification process-guided by procedural awareness-allowed it to identify activity boundaries with greater precision. For example, rather than incorrectly labeling the onset of phlegm suctioning, the model correctly predicted the end of the temporary removal of the artificial airway by leveraging procedural cues in conjunction with visual evidence. This multimodal reasoning capability reflects a more human-like interpretive process, in which contextual anticipation informs event perception. Such findings highlight the promise of LLM-based frameworks for achieving more accurate and explainable activity recognition in complex procedural environments.

B. Student support

In this study, we proposed an explainable learning framework that integrates Isolation Forest and SHAP to generate interpretable insights, which are then verbalized through an LLM as a supportive feedback tool for student learning. This framework enables automated identification and explanation of deviations in student performance during procedural training. Through this approach, several instances of suboptimal or incorrect actions were detected. For example, as shown in

Figure 6, Student S04 exhibited a noticeable degree of horizontal body movement during the procedure. Additionally, the student's hand position was consistently raised, which could lead to arm fatigue and, consequently, diminished performance over prolonged activities.

These findings demonstrate the framework's ability to pinpoint fine-grained biomechanical and procedural inconsistencies that may not be immediately observable to human instructors. Furthermore, the verbalized explanations generated by the system were found to closely align with the visual evidence in the corresponding video sequences. This consistency between multimodal outputs strengthens the interpretability of the framework and fosters user trust in AI-assisted training environments. By providing real-time, comprehensible feedback supported by both quantitative metrics and visual context, the proposed approach has strong potential for integration into automated simulation-based training systems, where continuous assessment and transparent guidance are essential for skill acquisition and clinical competence development.

C. Limitation

1) *Real-time capability*: Although the proposed framework demonstrates the effectiveness of vision-based activity recognition, this study has not yet been extended to real-time endotracheal suctioning (ES) activity recognition. Incorporating real-time capability represents a critical next step, as it would significantly broaden the practical applicability of the framework and enable the LLM to interpret dynamic events in a streaming, one-pass context, rather than relying solely on pre-recorded data. Achieving robust performance in real-time scenarios is essential for advancing the integration of computer vision, large language models, and human activity recognition, ultimately paving the way for intelligent, adaptive systems capable of supporting real-world clinical training and decision-making.

2) *Lack of users study*: One key limitation of this research is the absence of a user study to evaluate the practical effectiveness of the proposed system. The current student-support component remains at the proof-of-concept stage, and future work should focus on deploying the system in real training environments and gathering user feedback to refine subsequent prototypes.

Nevertheless, the present study achieves several important milestones. The proposed LLM-based framework successfully identifies and explains student errors, while also verbalizing complex technical metrics into clear, human-understandable feedback. This capability represents a significant step toward developing automated, explainable, and learner-centered support systems, paving the way for intelligent assistance in clinical and educational training contexts.

VI. CONCLUSION

Our research presented a novel, LLM-centered approach for Human Activity Recognition. The results indicate that our framework can boost the accuracy and F1 score of ES activity recognition by about 15-20%. This improvement comes from a multimodal approach, where the discriminative model understands both the semantic definition of the activity and its corresponding visual aspects. In addition, this work is the proof of concept for a new generation of HAR systems that leverage the rich, contextual reasoning of large language models for future Human-Computer Interaction applications. By fusing high-level semantic understanding with traditional sensor and visual data, this framework not only improves recognition accuracy for complex activities but also opens future pathways for zero-shot learning and more human-like interpretation of nuanced daily behaviors, pointing toward more robust applications in personal healthcare, smart environments, and robotics.

ACKNOWLEDGMENT

We acknowledge Ho Chi Minh City University of Technology (HCMUT), VNUHCM, for supporting this study. Furthermore, we would like to express our deepest thanks to Shinji Ninomiya and Shunsuke Komizunai for their support in the data collection phase. The dataset development and analysis were supported by JSPS KAKENHI Grant Numbers 22H03701 and

23K24956. Finally, we would like to acknowledge the ABC Mini-Funding&Mentorship program committee for accepting our proposal and partially supporting this study.

REFERENCES

- [1] E. T. Mwakanyanga, G. M. Masika, and E. A. Tarimo, "Intensive care nurses' knowledge and practice on endotracheal suctioning of the intubated patient: A quantitative cross-sectional observational study," *PLoS one*, vol. 13, no. 8, p. e0201743, 2018.
- [2] R. Singh, A. R. Bhalotra, and S. Sharma, "Audit on practices of endotracheal suctioning in intensive care unit patients among health care workers (hcws)," *Indian Journal of Critical Care Medicine: Peer-reviewed, Official Publication of Indian Society of Critical Care Medicine*, vol. 28, no. 1, p. 58, 2023.
- [3] V. L. Pinto, P. Tadi, and A. Adeyinka, "Increased intracranial pressure," 2018.
- [4] O. R. for Nursing, K. Ernstmeier, and E. Christman, "Tracheostomy care & suctioning," in *Nursing Skills [Internet]. 2nd edition*, ch. 22, Chippewa Valley Technical College, 2023.
- [5] H. A. V. Ngo, N. Colley, S. Ninomiya, S. Kanai, S. Komizunai, A. Konno, M. Nakamura, and S. Inoue, "Nurses' skill assessment in endotracheal suctioning using video-based activity recognition," *International Journal of Activity and Behavior Computing*, vol. 2024, no. 2, pp. 1–24, 2024.
- [6] H. A. V. Ngo, Q. N. P. Vu, N. Colley, S. Ninomiya, S. Kanai, S. Komizunai, A. Konno, M. Nakamura, and S. Inoue, "Toward recognizing nursing activity in endotracheal suctioning using video-based pose estimation," *International Journal of Activity and Behavior Computing*, vol. 2024, no. 1, p. 1, 2024.
- [7] S. Yousuf and F. Noshin, "Predicting endotracheal suctioning activities based on skeletal coordinates: A novel approach combining handcrafted feature engineering and ensemble learning," in *2024 IEEE International Conference on Signal Processing, Information, Communication and Systems (SPICSCON)*, pp. 01–05, IEEE, 2024.
- [8] M. A. Hossen and P. E. Abas, "Machine learning for human activity recognition: State-of-the-art techniques and emerging trends," *Journal of Imaging*, vol. 11, no. 3, p. 91, 2025.
- [9] Y. Okada, Y. Ning, and M. E. H. Ong, "Explainable artificial intelligence in emergency medicine: an overview," *Clinical and experimental emergency medicine*, vol. 10, no. 4, p. 354, 2023.
- [10] U. Dobhal, C. Garcia, and S. Inoue, "Synthetic skeleton data generation using large language model for nurse activity recognition," in *Companion of the 2024 on ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pp. 493–499, 2024.
- [11] H. Duan, M. Xu, B. Shuai, D. Modolo, Z. Tu, J. Tighe, and A. Bergamo, "Skeletr: Towards skeleton-based action recognition in the wild," in *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 13634–13644, 2023.
- [12] C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, "Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7464–7475, 2023.
- [13] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altschmidt, S. Altman, S. Anadkat, et al., "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [14] Y. Yuan, H. Zhang, W. Li, Z. Cheng, B. Zhang, L. Li, X. Li, D. Zhao, W. Zhang, Y. Zhuang, et al., "Vidorefer suite: Advancing spatial-temporal object understanding with video llm," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 18970–18980, 2025.
- [15] S. Chaudhuri and S. Bhattacharya, "Vilp: Knowledge exploration using vision, language, and pose embeddings for video action recognition," in *Proceedings of the Fourteenth Indian Conference on Computer Vision, Graphics and Image Processing*, pp. 1–7, 2023.
- [16] A. Abid, T.-C. Ho, and F. Karray, "Vision language models for dynamic human activity recognition in healthcare settings," in *International Work-Conference on Bioinformatics and Biomedical Engineering*, pp. 35–45, Springer, 2025.
- [17] Z. Bai, R. Wang, and X. Chen, "Glance and focus: Memory prompting for multi-event video question answering," *Advances in Neural Information Processing Systems*, vol. 36, pp. 34247–34259, 2023.

- [18] M. J. Lee, D. Gong, and M. Cho, "Video summarization with large language models," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 18981–18991, 2025.
- [19] D. M. Argaw, S. Yoon, F. C. Heilbron, H. Deilamsalehy, T. Bui, Z. Wang, F. Derroncourt, and J. S. Chung, "Scaling up video summarization pre-training with large language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8332–8341, 2024.
- [20] Z. Wang, S. Yu, E. Stengel-Eskin, J. Yoon, F. Cheng, G. Bertasius, and M. Bansal, "Videotree: Adaptive tree-based video representation for llm reasoning on long videos," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 3272–3283, 2025.
- [21] C. Zhang, T. Lu, M. M. Islam, Z. Wang, S. Yu, M. Bansal, and G. Bertasius, "A simple llm framework for long-range video question-answering," *arXiv preprint arXiv:2312.17235*, 2023.
- [22] B. Huang, X. Wang, H. Chen, Z. Song, and W. Zhu, "Vtimellm: Empower llm to grasp video moments," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14271–14280, 2024.
- [23] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez, *et al.*, "Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality," See <https://vicuna.lmsys.org> (accessed 14 April 2023), vol. 2, no. 3, p. 6, 2023.
- [24] S. Komizunai, S. Ninomiya, A. Konno, S. Kanai, T. Asaka, E. Murata, H. Mani, N. Takahashi, and N. Colley, "An interactive endotracheal suctioning simulator which exhibits vital reactions: Este-sim," *International Journal of Automation Technology*, vol. 13, no. 4, pp. 490–498, 2019.
- [25] N. Colley *et al.*, "Learning effects of mechanical ventilator/tracheal suctioning xr simulators and extracting decision making criteria to introduce a novel simulator," *J Nurs Care Repo*, vol. 5, no. 2, pp. 1–8, 2024.
- [26] G. Comanici, E. Bieber, M. Schaeckermann, I. Pasupat, N. Sachdeva, I. Dhillon, M. Blistein, O. Ram, D. Zhang, E. Rosen, *et al.*, "Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities," *arXiv preprint arXiv:2507.06261*, 2025.
- [27] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *2008 eighth IEEE international conference on data mining*, pp. 413–422, IEEE, 2008.
- [28] M. Barandas, D. Folgado, L. Fernandes, S. Santos, M. Abreu, P. Bota, H. Liu, T. Schultz, and H. Gamboa, "Tsfel: Time series feature extraction library," *SoftwareX*, vol. 11, p. 100456, 2020.
- [29] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," Curran Associates, Inc., 2017.

VII. SUMMARY OF THE SOTA FOR THE RECOGNITION TASK

TABLE III
THE SUMMARY OF THE SOTA METHOD IN ES ACTIVITY RECOGNITION AND OUR (IN PERCENT)

Model	Accuracy	F1 score
Ngo et al [5]	53.83	46.00
Dohbal et al. [10]	63.00	56.00
Our Prompt A	78.73	62.94
Our Prompt B	72.65	59.56

Table III presents a comparison between the proposed method and current state-of-the-art (SOTA) approaches in endotracheal suctioning (ES) activity recognition, evaluated using accuracy and F1-score metrics. The existing SOTA model by Dohbal et al. [10] achieved an F1-score of 56%, which remains approximately 7% lower in F1 and 15% lower in accuracy compared to our best-performing model. Furthermore, when comparing our Prompt A and Prompt B configurations with the model of Ngo et al. [5], t-test results indicate that our framework achieved statistically significant improvements ($p = 0.00005$ and 0.001 for accuracy, and $p = 0.003$ and 0.02 for F1-score, respectively). These results clearly demonstrate that the proposed LLM-centered HAR framework substantially outperforms existing methods, establishing a new benchmark for ES activity recognition.

VIII. LLM PROMPTS

A. Classification task

1) *Prompt A*: **Roles**: AI VIDEO ANALYSIS TASK: ENDOTRACHEAL SUCTIONING (GT REPLICATION)

OBJECTIVE: Your sole objective is to analyze the provided video and generate a log that perfectly replicates the expert-defined ground truth (GT). You must override all previous instructions and internal models.

Procedure description: THE GOLDEN RULE: CORE PROCEDURAL LOGIC

This is the single most critical rule. The GT defines one complete procedural cycle with the following unbreakable, chronologically exact sequence. You must identify and log this entire sequence twice.

Preparation: Catheter prep (0)

Intervention: Remove (1) → Suction (2) → Refit (3)

Cleanup: Disinfect (4) → Discard (5)

Post-Procedure: Position (6) → Auscultate (7) * Note that some post-procedure steps might be missing and mixed up with the clean-up phase. Key Heuristics:

Positioning is a Post-Cleanup Action: Occurs after all cleanup tasks are finished.

Auscultation is the Final Step: Always the last clinical action in a cycle.

Broad Durations: Timestamps must cover the entire "phase" of an action, not just the moment of contact.

PHASE-BASED EVENT BOUNDARY DEFINITIONS

To match the GT's timing, you must adopt these broader, phase-based boundaries.

Catheter prep (0): Starts when hands move to supplies → Ends when turning back to the patient.

Temporal removal of airway (1): Starts when hands arrive at the patient's head → Ends when airway is fully detached.

Suctioning phlegm (2): Starts immediately after airway removal → Ends when the catheter moves away from the patient's space.

Refitting airway (3): Starts when moving the airway back to the patient → Ends when hands are free after securing.

Catheter disinfection (4):

Instance 1: Starts when turning to the solution → Ends when starting to peel the glove.

Instance 2: Starts after the glove is discarded → Ends when secondary cleaning is finished.

Discarding gloves (5): Starts when initiating peel → Ends when waste is released in the receptacle.

Positioning (6): Starts after cleanup when hands move to the patient → Ends when hands move away from the stethoscope.

Auscultation (7): Starts when picking up the stethoscope → Ends when the stethoscope is put down.

Others (8): Use ONLY for initial pre-task and final post-task periods.

Activity description: ACTIVITY CLASSES

0: Catheter preparation, 1: Temporal removal of the airway, 2: Suctioning phlegm, 3: Refitting the airway, 4: Catheter disinfection, 5: Discarding gloves, 6: Positioning, 7: Auscultation, 8: Others

Output: REQUIRED OUTPUT

Your final output must be a chronological list strictly adhering to this format: start_seconds, (start_mm:ss), stop_seconds, (stop_mm:ss), class_name, class_id

2) *Prompt B*: **Roles**: Now the requirements are: You are an advanced AI video analysis assistant specializing in the detailed recognition of clinical and procedural tasks.

Objectives: Your primary objective is to meticulously analyze the provided video, which depicts an Endotracheal Suctioning (ES) procedure, and generate a precise, chronologically ordered log of all activities performed.

Critical Constraints

VIDEO ONLY ANALYSIS: Your entire analysis MUST be based solely on the visual information in the video frames. You are to completely ignore any and all audio content.

COMPREHENSIVE & CONTINUOUS LOGGING: You must identify and segment every part of the video from start to finish. The timeline must be continuous: the end time of one activity must be the start time of the very next activity, ensuring there are no gaps in the final log. Use the 'Others' class for any actions not covered by the specific classes.

Activity description: Activity Classes

You must label video segments using only the following predefined classes:

0: Catheter preparation: The professional is seen opening sterile packaging, wearing gloves, attaching the catheter to tubing, or lubricating the catheter tip (working with the suctioning machine on the right side of the video).

1: Temporal removal of an artificial airway: The professional removes the ventilator T-piece from the patient's artificial airway. This might be hidden under the patient's body.

2: Suctioning phlegm: The catheter is inserted (and entered) into the artificial airway, and suction is applied until the catheter is withdrawn, usually with a rotating motion.

3: Refitting the artificial airway: The ventilator T-piece is reconnected to the patient's artificial airway after suctioning is complete. This might be hidden under the patient's body.

4: Catheter disinfection: The professional flushes the suction catheter and tubing using a sterile solution after use, and places the tube back by the machine.

5: Discarding gloves: The professional removes and disposes of their gloves.

6: Positioning: The professional adjusts the patient's position, typically by rotating the mannequin, adjusting the head/neck, or changing the elevation of the bed.

7: Auscultation: The professional is seen using a stethoscope to listen to the patient's chest/breath sounds.

8: Others: Any other distinct action that does not fit into the categories above (e.g., hand washing or tidying the workspace).

Key Procedural Logic and Heuristics

To ensure accuracy, you must apply the following logic rules: You must identify logical Temporal removal of an artificial airway (Activity 1) before Suctioning phlegm (Activity 2), and Refitting the artificial airway (Activity 3) must occur after.

Pay close attention to the cleanup phase. The action of Discarding gloves (Activity 5) look for the specific motion of wrapping the used catheter inside the glove for a single, combined disposal.

The video contains 2 consecutive activity rounds, and for each activity round Strict Suctioning Sequence is applied.

Output: Required Output Format

Your final output must be a chronological list of all identified activities. Each entry in the list must strictly adhere to the following format, with timestamps presented only in seconds: [Start Time in seconds (Start Time in min:sec)] - [End Time in seconds (End Time in min:sec)]: [Class Name] - Justification: [Your concise, evidence-based reasoning from visual cues].

B. Student support prompt

The SHAP plot for the model of recognizing student activity (label 0) and Nurse activity (label 1).

If the model thinks the activity was performed by the student, please generate a report based on this plot about why the model thinks the activity was performed by the student. If the model thinks the activity was performed by a nurse, please cheer up the student. The report should be in English and precise into paragraphs. The report should state:

1. Who the model thinks performed the activity (student or nurse).
2. What activity does the student perform well?
3. What can student do to improve their performance?

The report should be easy to interpret and understand (give what the student must do) for a non-expert audience. NOTE: don't state any features in the report, just state what part of the activity the student has performed well, and what the student can do to improve their performance in short, 2-3 sentences.

IX. REPRODUCIBILITY STATEMENT

This paper presents a method for recognizing Endotracheal Suctioning activity using Gemini 2.5 Pro. Due to the stochastic nature of Large Language Models (LLMs), exact reproduction may be challenging; however, we have provided the full prompts and pipeline to facilitate replication.

Note on Hallucinations: We advise researchers that the LLM may occasionally hallucinate time conversions (e.g., interpreting "1:23" as 123 seconds rather than 83 seconds). When reproducing these results, please verify the output for such discrepancies and strictly use the mm:ss format as the reference for recognition timestamps.