

AURA: Development and Validation of an Augmented Unplanned Removal Alert System Using Synthetic ICU Videos

Junhyuk Seo^{1,2}, Hyeyoon Moon¹, Kyu-Hwan Jung¹, Namkee Oh³, Taerim Kim^{1,4,*}

¹Samsung Advanced Institute for Health Sciences & Technology, Sungkyunkwan University

²Department of Nursing, Samsung Medical Center

³Department of Surgery, Samsung Medical Center, Sungkyunkwan University School of Medicine

⁴Department of Emergency Medicine, Samsung Medical Center, Sungkyunkwan University School of Medicine
{junhyuk.sam, moon012427, kyuhwanjung, ngnyou, taerim.j.kim}@gmail.com

Abstract

Unplanned extubation (UE)—the unintended removal of an airway tube—remains a critical patient safety concern in intensive care units (ICUs), often leading to severe complications or death. Real-time UE detection has been limited, largely due to the ethical and privacy challenges of obtaining annotated ICU video data. We propose Augmented Unplanned Removal Alert (AURA), a vision-based risk detection system developed and validated entirely on a fully synthetic video dataset. By leveraging text-to-video diffusion, we generated diverse and clinically realistic ICU scenarios capturing a range of patient behaviors and care contexts. The system applies pose estimation to identify two high-risk movement patterns: collision, defined as hand entry into spatial zones near airway tubes, and agitation, quantified by the velocity of tracked anatomical keypoints. Expert assessments confirmed the realism of the synthetic data, and performance evaluations showed high accuracy for collision detection and moderate performance for agitation recognition. This work demonstrates a novel pathway for developing privacy-preserving, reproducible patient safety monitoring systems with potential for deployment in intensive care settings.

Code — <https://github.com/seo-see/AURA>

Datasets — <https://doi.org/10.5281/zenodo.17577177>

Introduction

Mechanical ventilation via artificial airway is a lifesaving intervention for critically ill patients in intensive care units (ICUs) (Chao et al. 2017). Paradoxically, one of the most critical safety concerns in ICUs is unplanned extubation (UE)—the unintended removal of the airway tube (Christie, Dethlefsen, and Cane 1996). A systematic review reported that approximately 6.7% of mechanically ventilated patients experience UE, and over 80% of these cases are self removal (Li, Sun, and Xu 2022). Such events are associated with prolonged ICU stays, increased healthcare costs, higher rates of infection, and elevated mortality (da Silva and Fonseca 2012). In response, healthcare providers have implemented various preventive strategies, including physical restraints, standardized sedation protocols, and routine delir-

ium screening (Chang, Wang, and Chao 2008; Jarachovic et al. 2011; Alaterre et al. 2023).

Several studies have identified nurse unavailability at the bedside as the dominant factor contributing to UE, accounting for nearly 90% of incidents (Curry et al. 2008; Selvan et al. 2014), and insufficient nurse staffing has been consistently linked to higher UE rates (Penoyer 2010). These findings indicate that continuous bedside oversight is critical for UE prevention but is difficult to sustain in real-world ICUs. The COVID-19 pandemic further exacerbated this challenge, as staffing shortages and strict isolation protocols hindered patient monitoring, coinciding with an increased UE incidence (Berkow and Kanowitz 2020; Taylor, Reynolds, and Jones 2021). In these low-resource and high-burden settings, video monitoring has been proposed as a potential solution (Taylor, Reynolds, and Jones 2021).

However, conventional video monitoring relies on human observers and operates passively (Cournan, Fusco-Gessick, and Wright 2016; Abbe and O’Keeffe 2021). Although active patient monitoring using wearable sensors and Internet of Things (IoT) devices has been explored for patient safety (Chiang and Liu 2019; Joseph et al. 2024), these methods impose additional maintenance burdens on clinical staff. More recently, computer vision approaches have shown promise for non-contact, sensor-free patient monitoring. While fall detection and risk assessment have been widely studied to improve patient safety (Gharghan and Hashim 2024; Gabriel et al. 2025), UE-specific applications remain limited, largely due to the ethical and privacy concerns of ICU surveillance video (Chen et al. 2024).

To address these gaps, we propose Augmented Unplanned Removal Alert (AURA), a system that leverages synthetic ICU video data to detect UE risk. By using visually realistic yet privacy-preserving synthetic footage, AURA enables system development and expert evaluation without real ICU recordings—allowing refinement and failure mode analysis. The system relies on pose estimation of generalizable anatomical cues (e.g., mouth–hand proximity) instead of hand-crafted features, enabling future extension to other critical devices such as central venous catheters. At its current stage, AURA functions as a pre-deployment tool, demonstrating how synthetic video can enable privacy-preserving, reproducible development of real-time safety monitoring systems for future clinical integration.

*Corresponding author.

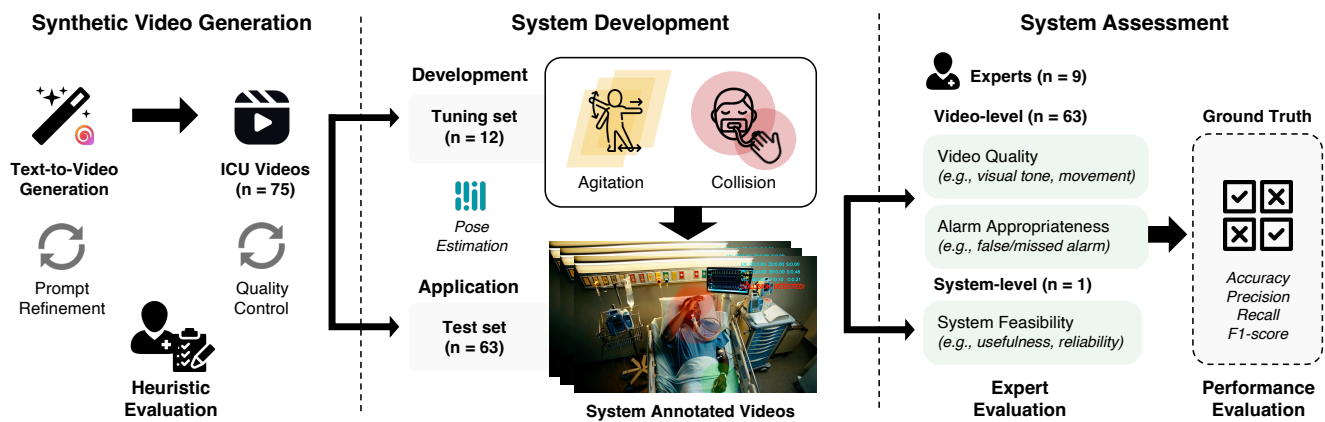


Figure 1: Overview of the AURA development pipeline. The framework consists of three stages: (1) Synthetic video generation: 75 ICU videos generated using text-to-video model and refined through heuristic evaluation; (2) System development: Pose estimation-based detection of collision and agitation behaviors, with model tuning (n=12) and testing (n=63); (3) System assessment: Nine experts conducted video- and system-level evaluation, establishing ground truth for subsequent performance evaluation.

Related Work

Data-Driven Approaches

Several studies have integrated machine learning (ML) models into clinical workflows to mitigate the risk of unplanned tube or catheter removal. Hur et al. (2021) developed and validated a prediction model for UE using electronic health record (EHR) data. Random forest (RF) was reported as the best-performing model, with an area under the curve (AUC) of 0.79 and a sensitivity of 0.95. Similarly, Zhang et al. (2023) proposed a support vector machine (SVM) model for central catheter unplanned removal, achieving an AUC of 0.88 and a sensitivity of 0.89. These ML models, however, primarily function as early-warning risk scoring tools based on historical features from EHR. They lack real-time awareness and cannot capture the physical movements that directly precede UE, leaving a gap between predictive warning and actionable patient monitoring.

Wearable Sensors

Wearable sensing has also been explored for patient safety monitoring in ICUs. Chiang and Liu (2019) developed a wearable infrared sensor using complementary metal-oxide-semiconductor (CMOS) technology to detect hand movements toward airway tubes and transmit proximity data wirelessly to nursing stations. Likewise, Joseph et al. (2024) applied accelerometer-based sensors for fall prevention, a conceptually similar safety task, by continuously capturing motion patterns to identify high-risk behaviors. While feasible, these wearable solutions require patient compliance and device maintenance, which may disrupt workflow and limiting large-scale deployment.

Computer Vision Approaches

Computer vision applications provide a non-contact alternative by directly observing patient behavior. Chen et al.

(2024) proposed a hand-crafted spatio-temporal feature framework that combined corner detection, trajectory distance, and wavelet transforms to characterize motion patterns. Among various models, SVM model achieved an accuracy of 0.85 for binary classification (UE tendency vs. no UE tendency) and 0.62 for three-class classification (including no movement). Although this approach demonstrated competitive performance, it required labor-intensive feature engineering and relied on limited, privacy-sensitive ICU video data, restricting both scalability and reproducibility.

Our Contributions

In contrast to previous approaches that rely on structured EHR data, wearable sensors, or real ICU surveillance footage, our proposed system introduces a fully synthetic, vision-based framework for real-time UE risk detection. We leverage pose estimation on synthetic videos to detect high-risk behaviors without the need for physical sensors, hand-crafted features. This enables scalable, reproducible real-time safety monitoring while preserving patient privacy.

Our system is also distinct in that it incorporates structured clinical expert evaluations, an underexplored aspect in prior UE risk detection research. Through expert assessments of system reliability and clinical feasibility, we establish a clinically grounded validation framework that enhances the interpretability and trustworthiness of the system prior to deployment.

Data and Methods

This section describes the development and validation of our synthetic ICU video dataset and the proposed vision-based system for UE risk detection. The process comprises three main stages: (1) Synthetic video generation, (2) System development, and (3) System assessment. Figure 1 provides an overview of this workflow.

Synthetic Video Generation and Dataset

The AURA dataset comprises 75 synthetic ICU videos depicting intubated patients. Each video is 6 seconds long, with a frame rate of 25 frame per second (fps) and a resolution of 1280×720 pixels. The videos were generated using Hailuo AI’s T2V-01-Director model (Hailuo AI 2024). This model was selected for its ability to produce static, surveillance-like scenes, while other tested models tended to generate continuously shifting, cinematic camera movements that were unsuitable for this study’s objective.

Video generation prompts were carefully designed and validated by an experienced ICU nurse. A structured prompt template was used to systematically create diverse scenarios varying by patient gendered appearance and racial appearance, behavior (e.g., calm, agitated, lifting hands, pulling tube), and the presence of medical staff in the scene. To enhance visual diversity, we explicitly included Hispanic/Latino as a distinct appearance category. We limited our dataset to 75 videos to balance diversity with practical constraints while ensuring adequate coverage of key variations and keeping the evaluation feasible for experts.

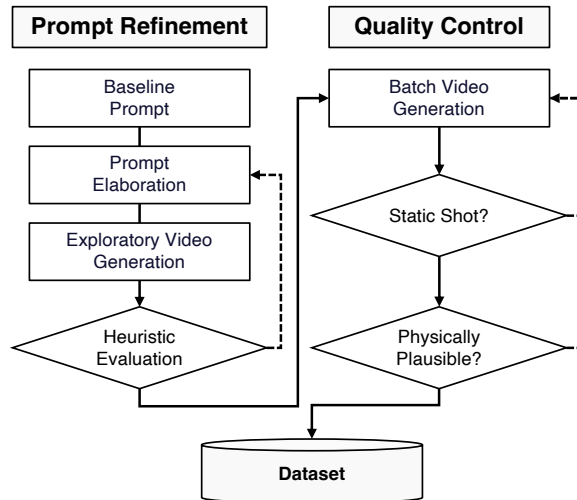


Figure 2: Overview of the AURA synthetic video dataset generation. Prompts are refined via exploratory generation and heuristic expert evaluation, then used for batch video generation. Outputs pass two screening stages and only accepted videos form the final dataset (solid = pass; dashed = feedback/fail). This integrated human–AI workflow ensures realism and feasibility while preserving reproducibility.

The generation process involved iterative prompt refinement through heuristic evaluation, including exploratory generation and quality control, to ensure realistic and contextually accurate ICU scenes (Figure 2). Prompts were tuned through small-scale generations, and outputs exhibiting cinematic shots or physically implausible patient movements were excluded. This process maintained both visual fidelity and reproducibility, providing a robust foundation for subsequent video- and system-level validation.

Each video required approximately two minutes to generate, and about 20% of the outputs failed quality control

and were re-generated. Importantly, no real-world resources (such as photographs or videos) were used as inputs during generation. The complete set of synthetic videos, along with representative prompts, is publicly available on Zenodo and a summary of the final dataset is shown on Table 1. Notably, the dataset’s gendered appearance distribution (56% male, 44% female) closely mirrors that reported among ICU patients worldwide (Merdji et al. 2023).

Attribute	Count (n)	Proportion (%)
Gendered appearance		
Male	42	56.0
Female	33	44.0
Racial appearance		
Black	21	28.0
Asian	20	26.7
White	20	26.7
Hispanic/Latino	14	18.7
Behavior		
Movement	49	65.3
Calm	26	34.7
Camera angle		
Vertical	61	81.3
Diagonal	7	9.3
Horizontal	7	9.3
Full body visibility		
No	66	88.0
Yes	9	12.0
Medical staff		
None	64	85.3
Interaction	8	10.7
Presence	3	4.0

Table 1: Descriptive statistics of the full synthetic ICU video dataset ($n = 75$). All appearance attributes reflect visually generated characteristics.

System Development

The purpose of this study is to detect risk of unplanned airway removal in ICU patients. We implemented a pose estimation model to enable future extensions beyond airway monitoring, including detection of risks related to other critical devices (e.g., central venous catheters and chest tubes). All thresholds were fixed by clinical expert after optimization on a 12-video tuning set. We then validated on 63 unseen videos to avoid leakage.

Multiple studies have identified agitation as the most significant patient-related risk factor for UE (Hur et al. 2021). Agitation, a key clinical indicator of UE risk, is formally captured by the Richmond Agitation-Sedation Scale (RASS) (Sessler et al. 2002). According to the RASS, agitation ranges from restless and anxious behavior to overtly combative actions. A very agitated state often includes direct risk behaviors such as “pulling on or removing tubes,”



Figure 3: Examples of overlays for collision (left) and agitation (right) in synthetic intensive care unit videos. Colored “auras” around hands and mouth indicate risk zones (green = normal; red = collision). Numeric overlays show hand proximity (LH, RH), movement velocity (VEL).

whereas moderately agitated states manifest as anxiousness, restlessness, and frequent non-purposeful limb movements.

To capture both immediate and early risks, we developed two behavior detection mechanisms: **collision**, which indicates purposeful hand movements toward the airway tube and requires immediate intervention, and **agitation**, a broader behavioral state that may predispose the patient to collision events. This distinction also reflects practical differences in detection: collision is spatially well-defined and more reliably detectable, while agitation is inherently more abstract and challenging to capture (Figure 3).

The **collision** detection algorithm integrates 2D overlap analysis and 3D proximity measurement to address the limitations of 3D depth estimation from monocular cameras. While 3D coordinates provide depth information, they are often noisy and unreliable due to variability of camera angle and full body invisibility. Therefore, we complement 3D proximity with 2D overlap analysis, which yields more stable and robust measurements in the image plane.

The 2D overlap score between mouth and hand auras is calculated as:

$$\text{overlap_score} = \max\left(0, 1 - \frac{d_{2D}}{r_h + r_m}\right)$$

where $d_{2D} = \sqrt{(x_h - x_m)^2 + (y_h - y_m)^2}$ is the 2D distance between hand and mouth coordinates, and r_h, r_m are the hand and mouth aura radii respectively.

The 3D proximity score is computed as:

$$\text{proximity_score} = \max\left(0, 1 - \frac{d_{3D}}{\tau_{base}}\right)$$

where $d_{3D} = \sqrt{(x_h - x_m)^2 + (y_h - y_m)^2 + (z_h - z_m)^2}$ is the 3D distance and τ_{base} is the distance threshold.

The final collision score combines both metrics with weighted averaging:

$$\text{score} = \alpha \cdot \text{overlap_score} + \beta \cdot \text{proximity_score}$$

where α and β denote the respective weights. A collision is triggered when the final score exceeds the threshold τ_{score} . To mitigate false alarms caused by anchor instability, we apply a persistence-based mechanism that confirms a collision only if the risk state lasts longer than a minimum duration ($\tau_{duration}$). In addition, keypoints with visibility scores below τ_{valid} are excluded from the analysis to reduce noise from low-confidence detections.

For **agitation** detection, we quantify movement intensity based on the 3D velocity of each keypoint i , defined as v_i , where d_i represents the 3D displacement between consecutive frames:

$$v_i = \frac{d_i}{\Delta t},$$

$$d_i = \sqrt{(x_i^t - x_i^{t-1})^2 + (y_i^t - y_i^{t-1})^2 + (z_i^t - z_i^{t-1})^2}.$$

Only keypoints with visibility $\geq \tau_{valid}$ are included to ensure reliability in motion estimation. For a temporal window of w frames, we derive three motion statistics:

$$\text{mean_velocity} = \frac{1}{|V_{valid}|} \sum_{v \in V_{valid}} v,$$

$$\text{peak_velocity} = \max_{v \in V_{valid}} v,$$

$$\text{cumulative_velocity} = \sum_{v \in V_{valid}} v.$$

Agitation is detected when any of the following conditions are met:

$$\text{is_agitation} = \begin{cases} \text{True} & \text{if (mean_velocity} > \tau_{speed}) \\ & \text{or (peak_velocity} > \tau_{speed}) \\ & \text{or (cumulative_velocity} > \tau_{speed} \cdot w) \\ \text{False} & \text{otherwise.} \end{cases}$$

A key feature of AURA is the visualization of augmented spatial zones for collision detection. Circles surrounding the patient’s hands and mouth—like an aura—appear green during normal status and change to red when a collision

is detected. These overlays were purposefully designed to enhance clinical interpretability and situational awareness, while they do not affect the underlying detection logic.

All system parameters were optimized through iterative clinical expert evaluation with tuning set of 12 videos. The expert reviewed system performance across various scenarios and adjusted parameters to achieve optimal balance between detection sensitivity and false alarm reduction (Table 2).

Parameter	Value	Description
τ_{base}	0.3	3D distance threshold
r_m	150 pixels	Mouth Aura radius
r_h	100 pixels	Hand Aura radius
α	0.7	2D overlap weight
β	0.3	3D proximity weight
τ_{score}	0.3	Final score threshold
$\tau_{duration}$	0.3s	Persistence duration
τ_{speed}	0.18	Agitation speed threshold
τ_{valid}	0.7	Keypoint visibility threshold
w	5 frames	Analysis window size

Table 2: Parameter settings for the AURA system

After prototype development, the AURA system was applied to the test set of 63 unseen videos. Each video was processed through the optimized detection pipeline, generating collision and agitation alerts based on clinically calibrated thresholds. As the primary objective was feasibility validation using pre-recorded data, real-time streaming integration was not implemented at this stage.

Pose estimation was performed using MediaPipe, selected for its cross-platform efficiency and low-latency stream processing capabilities (Lugaresi et al. 2019). All experiments were conducted with Python 3.9 and OpenCV on macOS (Apple M3 Max), without GPU acceleration. The source code is publicly available on our GitHub repository.

System Assessment

We designed a single integrated validation framework that simultaneously assessed the realism of the synthetic ICU videos and the performance of the AURA system. A structured web-based questionnaire, co-developed by clinical and informatics experts, holistically evaluated three domains: (1) the quality of the synthetic videos, (2) the clinical appropriateness of system-generated alarms, and (3) the feasibility of integrating the system into real-world ICU workflows.

Nine ICU nurses independently evaluated 63 annotated videos. The evaluators were evenly distributed across three experience levels (early-career: 1–3 years, mid-career: 3–5 years, and senior-level: over 5 years) to ensure diverse clinical perspectives. For each video, they rated four dimensions of visual realism (visual tone, background setting, patient behavior, and clinical plausibility) as well as alarm appropriateness using a 5-point Likert scale (1 = strongly disagree, 5 = strongly agree). To assess the reliability of these expert ratings, intraclass correlation coefficients (ICCs) were calculated for alarm appropriateness scores.

When an alarm was deemed clinically inappropriate, evaluators selected one of four predefined categories: false alarm, missed alarm, premature alarm, or delayed alarm. These annotations provided both qualitative insights and consensus-based reference labels for computing performance metrics (accuracy, precision, recall, and F1-score). 95% confidence intervals (CI) were estimated using bootstrap resampling. Notably, alarms classified as premature or delayed were still counted as system-triggered alarms during evaluation, reflecting a conservative labeling strategy that prioritizes sensitivity over exact temporal alignment.

After video-level evaluation, system-level feasibility was also assessed with the same Likert scale across three criteria (clinical usefulness, reliability, and recommendation to others).

Robustness Analysis

Two complementary analyses were conducted to evaluate system robustness with respect to both spatial scale and parameter perturbation.

Scale Robustness. As the system initially relied on a fixed pixel aura, variations in video size, camera zoom, viewing angle, and patient body size could potentially influence model performance. To examine scale robustness, we tested a relative pixel aura by normalizing all distance measurements to each subject’s head and hand dimensions. The normalized body sizes were computed as follows:

$$\begin{aligned} \text{head_size} = & \max(\text{dist}(L_{\text{ear}}, R_{\text{ear}}), \\ & 2 \times \text{dist}(L_{\text{eyebrow}}, L_{\text{mouth}}), \\ & 2 \times \text{dist}(R_{\text{eyebrow}}, R_{\text{mouth}})) \end{aligned}$$

$$\begin{aligned} \text{hand_size} = & \max(\max(\text{dist}(L_{\text{pinky}}, L_{\text{thumb}}), \text{dist}(L_{\text{wrist}}, L_{\text{index}})), \\ & \max(\text{dist}(R_{\text{pinky}}, R_{\text{thumb}}), \text{dist}(R_{\text{wrist}}, R_{\text{index}}))) \end{aligned}$$

Let λ denote the scaling coefficient for the aura radii. The final aura radii were computed as:

$$\begin{aligned} \text{head_aura_radius} &= \lambda \cdot \text{head_size}, \\ \text{hand_aura_radius} &= \lambda \cdot \text{hand_size}. \end{aligned}$$

In our implementation, λ was empirically set to 2.0 to balance sensitivity and stability in the tuning set. Additionally, we tested relative mode for down-scaled video (from 1280×720 to 854×480) to demonstrate the system robustness in low resolution video.

Parameter Robustness. We conducted a concise, 3-fold cross-validation–style sensitivity analysis. In each fold, 21 videos were used as a tuning set and 42 as a held-out validation set. Key hyperparameters were perturbed within $\pm 10\%$ of their expert-defined baselines, and a 3^3 grid search (27 configurations) was performed over the agitation speed threshold (τ_{speed}), visibility threshold (τ_{valid}), and radius scaling factor (s_r), which scales the fixed pixel radius of the aura regions as $r_{\text{final}} = s_r \times r_{\text{base}}$.

The selection objective on the tuning set was the combined F1 score (mean of collision and agitation-F1), which was then evaluated on the held-out set. This procedure was repeated over three non-overlapping validation splits.

Domain	Definition	Mean \pm SD	Range
Visual tone	Visual tone (brightness/contrast) of the video is similar to clinical setting	4.4 \pm 0.1	4.0–4.7
Background setting	Background of the video is similar to clinical setting	4.3 \pm 0.2	3.3–4.6
Patient behavior	Behavior of the patient in the video seems natural	4.6 \pm 0.2	3.9–4.8
Clinical plausibility	The event presented in the video is plausible in the clinical context	4.6 \pm 0.2	4.2–4.9

Table 3: Evaluation framework for synthetic ICU video dataset (5-point Likert, with range)

Results

Expert Evaluation

In the video quality assessment, the synthetic ICU videos were perceived as highly realistic and clinically valid. All four evaluation domains (visual tone, background setting, patient behavior, and clinical plausibility) received mean Likert scores ≥ 4.3 (range: 3.3–4.9), with patient behavior and clinical plausibility showing the highest ratings at 4.6 (Table 3).

For alarm appropriateness, mean video-wise ratings with standard deviation (SD) indicated high appropriateness for collision alarms (4.3 \pm 0.6) and moderate appropriateness for agitation alarms (3.8 \pm 0.6). Inter-rater reliability was excellent for both alarm types, with ICC(3,k) values confirming consistent expert agreement (Table 4).

Domain	ICC(3,k)	95% CI	F	p-value
Agitation	0.95	[0.89, 0.99]	19.62	<0.001
Collision	0.95	[0.89, 0.99]	19.42	<0.001

Table 4: Inter-rater reliability for alarm appropriateness evaluation by nine experts, calculated using the intraclass correlation coefficient ICC(3,k). Degrees of freedom for the F-statistic: $df_1 = 8$, $df_2 = 496$.

System-level feasibility scores showed medium-to-high acceptance, with mean ratings of 3.9 \pm 0.9 for clinical usefulness, 3.2 \pm 1.0 for reliability, and 3.9 \pm 0.8 for recommendation to others.

Performance Evaluation

The collision detection module achieved near-perfect performance, with F1-score of 0.98 and recall of 1.00. In contrast, the agitation detection module demonstrated moderate performance (F1-score: 0.78), reflecting the complexity of capturing agitation behaviors (Table 5).

Metric	Collision Detection	Agitation Detection
Accuracy	0.98 [0.93, 1.00]	0.86 [0.76, 0.95]
Precision	0.96 [0.81, 0.99]	0.89 [0.74, 1.00]
Recall	1.00 [0.86, 1.00]	0.70 [0.51, 0.88]
F1-score	0.98 [0.93, 1.00]	0.78 [0.62, 1.00]

Table 5: Performance metrics of AURA with 95% confidence intervals.

Robustness Analysis

The relative aura radii achieved performance comparable to the fixed-radius baseline, with similar F1 and recall scores for both behaviors (collision: F1 = 0.94, recall = 1.00; agitation: F1 = 0.78, recall = 0.70). When tested on down-scaled videos (854 \times 480), collision detection remained highly robust (F1 = 0.92, recall = 1.00) while agitation detection showed moderate degradation (F1 = 0.65).

Across the three validation folds, collision detection maintained stable performance (F1 = 0.94–1.00, recall = 1.00), whereas agitation detection showed moderate variability (F1 = 0.69–0.83, recall = 0.60–0.75). The best-performing configuration was consistent across folds, with $\tau_{\text{speed}} = 0.18$, $\tau_{\text{valid}} = 0.63$, and $s_r = 0.9$. The overall cross-fold deviation in combined F1 was within 0.05, demonstrating stable performance under $\pm 10\%$ parameter perturbations.

Discussion

The proposed system demonstrates the feasibility and ethical viability of developing a real-time early warning system for UE using only synthetic videos. By integrating pose estimation with realistic synthetic data, AURA bridges the gap between traditional risk-scoring models and continuous behavioral monitoring. This represents a practical and ethical alternative to real ICU video collection, overcoming privacy and logistical barriers in high-acuity settings.

Synthetic ICU videos proved to be a viable resource for both system development and expert validation. High evaluation scores (all ≥ 4.3) across realism domains indicate that synthetic videos can replicate key visual cues necessary for risk assessment. These results validate the use of synthetic datasets for scalable and privacy-preserving development of vision-based patient monitoring tools.

The reliability of alarm appropriateness assessments, which served as the reference standard for performance evaluation, was supported by strong inter-rater agreement. The ICC(3,k) exceeded 0.95 for both collision and agitation ratings, demonstrating the robustness and consistency of the evaluation framework. In addition to validating the system’s performance, this work contributes a novel expert-annotated dataset. Given the inherently ambiguous and context-dependent nature of agitation, these annotations capture subtle clinical judgments that cannot be fully represented through rule-based definitions. This resource provides a valuable foundation for future research on behavioral monitoring and UE risk stratification.

Subsequent performance evaluation revealed a clear distinction in system accuracy across behavior types. Collision

detection achieved near-perfect results (F1: 0.98; Recall: 1.00), benefiting from its spatially well-defined characteristics and the high contrast between normal and high-risk hand trajectories. In contrast, agitation detection yielded moderate performance (F1: 0.78; Recall: 0.70), reflecting the broader variability and subtlety of agitation-related movements. Future work should focus on improving quantification and categorization of agitation using advanced temporal modeling and expanded behavior taxonomies.

False alarms were primarily attributed to two sources: anchor instability and staff interference. Anchor jumps (momentary misplacement of pose keypoints near critical zones) occasionally triggered false collision alerts. Similarly, interaction between patient and medical staff in the frame led to misidentified limb movements. These issues point to limitations of the current pose estimation backend (MediaPipe), which lacks temporal smoothing and multi-person detection. Incorporating temporal filtering and gating mechanisms, as proposed in recent work (Chen et al. 2024), may reduce false positives and improve clinical reliability.

Beyond quantitative performance, expert feedback also revealed the system’s practical limitations, particularly regarding alert precision. The overall clinical reliability was rated moderate (3.2/5), reflecting the high standards of accuracy expected by ICU professionals. These findings emphasize the importance of aligning algorithmic detection thresholds with clinical expectations. Accordingly, this study introduces one of the first clinician-centered evaluation frameworks for UE detection, providing a replicable foundation for validating safety-critical monitoring systems.

Despite these practical limitations, AURA demonstrated stable detection performance when tested with relative aura radii, under $\pm 10\%$ parameter perturbations, and across varying video resolutions. Collision detection remained highly accurate (F1 = 0.92) even at reduced resolution (854×480 , 67% of original), indicating strong robustness to video quality degradation. Although agitation detection exhibited a moderate performance decline at lower resolutions, this is unlikely to pose significant constraints given the widespread adoption of high-definition cameras (1280×720) in modern ICU surveillance systems. These findings demonstrate that the system’s underlying detection logic is stable to moderate variations in both algorithmic parameters and input quality, supporting potential deployment in real-world ICUs.

Although this study evaluated system outputs on pre-recorded video, the same detection logic can be deployed in real time by subscribing to existing Closed-Circuit Television (CCTV) streams, requiring no new sensors or camera replacements. Deploying AURA in real-world ICUs would enable continuous UE risk monitoring, bringing this approach closer to clinical deployment.

To support such deployment, future iterations of AURA should incorporate adaptive thresholding, temporal stabilization, and robust multi-person tracking. Together with real-time evaluation in actual ICU workflows, these enhancements will be essential for ensuring safe, trustworthy, and clinically reliable vision assistance in critical care settings.

Limitations

Our fully synthetic dataset may lack rare postures, extreme agitation, or uncommon device configurations, limiting generalizability to complex ICU scenarios. Visible phenotype categories serve solely as visual proxies and do not capture the full diversity of real-world patients.

Despite high scores for realism and plausibility, the synthetic videos may differ from real ICU footage in lighting, saturation, or camera perspective. To improve robustness, future work should incorporate diverse visual augmentations. In addition, overlays (“aura” visualizations) were visible during expert evaluations and may have influenced perceptions of video quality and event plausibility. We retained them to reduce evaluators’ cognitive load during rapid assessments, but will conduct overlay-free evaluations in future studies to better estimate potential bias.

Conclusion

We introduced AURA, a pose-estimation-based early warning system for UE, developed and validated entirely on a fully synthetic ICU video dataset. Without relying on real ICU footage or patient-worn sensors, AURA demonstrates the feasibility of scalable, privacy-preserving behavioral monitoring in critical care.

The synthetic videos achieved high realism across all evaluation domains (mean Likert ≥ 4.3), and nine ICU nurses reached excellent agreement (ICC > 0.95) when assessing system-generated alarms. Based on this rigorously validated ground truth, AURA achieved near-perfect collision detection (F1 = 0.98) and respectable agitation recognition (F1 = 0.78). Subsequent robustness analysis confirmed stable performance of the system with moderate variations in parameters and relative mode settings.

Our findings establish synthetic video as a valuable and viable resource for developing safety-critical AI in privacy-sensitive healthcare settings. Beyond UE monitoring, AURA’s reproducible pipeline provides a blueprint for ethical AI development that can be extended to other critical devices and patient safety applications. This work demonstrates that privacy-preserving approaches can deliver clinically meaningful, deployable monitoring solutions for intensive care environments worldwide.

Ethics Statement

This study was reviewed and approved by the Institutional Review Board (IRB) of Samsung Medical Center (IRB No. SMC 2025-06-135-001) and conducted in accordance with all relevant ethical standards and institutional regulations.

No real patient data or footage were used in the synthetic video generation process. All participating clinical experts provided written informed consent, which covered the study purpose, the collection and analysis of their evaluation responses, and the potential public release of de-identified, aggregated data (e.g., Likert-scale scores) for secondary academic use.

Acknowledgments

This research was supported by a grant of the Korea Health Technology R&D Project through the Korea Health Industry Development Institute (KHIDI), funded by the Ministry of Health & Welfare, Republic of Korea (grant number : RS-2025-02223382).

References

- Abbe, J. R.; and O’Keeffe, C. 2021. Continuous video monitoring: Implementation strategies for safe patient care and identified best practices. *Journal of nursing care quality*, 36(2): 137–142.
- Alaterre, C.; Fazilleau, C.; Cayot-Constantin, S.; Chanques, G.; Kacer, S.; Constantin, J.-M.; and James, A. 2023. Monitoring delirium in the intensive care unit: Diagnostic accuracy of the CAM-ICU tool when performed by certified nursing assistants—A prospective multicenter study. *Intensive and Critical Care Nursing*, 79: 103487.
- Berkow, L.; and Kanowitz, A. 2020. COVID-19 putting patients at risk of unplanned extubation and airway providers at increased risk of contamination. *Anesthesia & Analgesia*, 131(1): e41–e43.
- Chang, L.-Y.; Wang, K.-W. K.; and Chao, Y.-F. 2008. Influence of physical restraint on unplanned extubation of adult intensive care patients: a case-control study. *American Journal of Critical Care*, 17(5): 408–415.
- Chao, C.-M.; Lai, C.-C.; Chan, K.-S.; Cheng, K.-C.; Ho, C.-H.; Chen, C.-M.; and Chou, W. 2017. Multidisciplinary interventions and continuous quality improvement to reduce unplanned extubation in adult intensive care units: A 15-year experience. *Medicine*, 96(27): e6877.
- Chen, Y.; Wang, L.; Wang, G.; Yang, S.; Wang, Y.; Xiang, M.; Zhang, X.; Chen, H.; Hu, D.; and Cheng, H. 2024. Spatio-temporal features for fast early warning of unplanned self-extubation in icu. *Engineering Applications of Artificial Intelligence*, 127: 107294.
- Chiang, C.-T.; and Liu, K.-Y. 2019. A cmos wearable infrared light intensity digital converter for monitoring unplanned self-extubation of patients. *IEEE Sensors Journal*, 19(15): 6430–6436.
- Christie, J. M.; Dethlefsen, M.; and Cane, R. D. 1996. Unplanned endotracheal extubation in the intensive care unit. *Journal of clinical anesthesia*, 8(4): 289–293.
- Cournan, M.; Fusco-Gessick, B.; and Wright, L. 2016. Improving patient safety through video monitoring. *Rehabilitation Nursing*.
- Curry, K.; Cobb, S.; Kutash, M.; and Diggs, C. 2008. Characteristics associated with unplanned extubations in a surgical intensive care unit. *American Journal of Critical Care*, 17(1): 45–51.
- da Silva, P. S. L.; and Fonseca, M. C. M. 2012. Unplanned endotracheal extubations in the intensive care unit: systematic review, critical appraisal, and evidence-based recommendations. *Anesthesia & Analgesia*, 114(5): 1003–1014.
- Gabriel, P.; Rehani, P.; Troy, T.; Wyatt, T.; Choma, M.; and Singh, N. 2025. Continuous patient monitoring with AI: real-time analysis of video in hospital care settings. *Frontiers in Imaging*, 4: 1547166.
- Gharghan, S. K.; and Hashim, H. A. 2024. A comprehensive review of elderly fall detection using wireless communication and artificial intelligence techniques. *Measurement*, 226: 114186.
- Hailuo AI. 2024. Hailuo AI. <https://hailuoai.video/landing-page>. Accessed: 2025-05-05.
- Hur, S.; Min, J. Y.; Yoo, J.; Kim, K.; Chung, C. R.; Dykes, P. C.; and Cha, W. C. 2021. Development and validation of unplanned extubation prediction models using intensive care unit data: Retrospective, comparative, machine learning study. *Journal of medical Internet research*, 23(8): e23508.
- Jarachovic, M.; Mason, M.; Kerber, K.; and McNett, M. 2011. The role of standardized protocols in unplanned extubations in a medical intensive care unit. *American journal of critical care*, 20(4): 304–312.
- Joseph, P.; Ali, H.; Matthew, D.; Thomas, A.; Jose, R.; Mayer, J.; Bekbolatova, M.; Devine, T.; and Toma, M. 2024. Regressive Machine Learning for Real-Time Monitoring of Bed-Based Patients. *Applied Sciences*, 14(21): 9978.
- Li, P.; Sun, Z.; and Xu, J. 2022. Unplanned extubation among critically ill adults: A systematic review and meta-analysis. *Intensive and Critical Care Nursing*, 70: 103219.
- Lugaresi, C.; Tang, J.; Nash, H.; McClanahan, C.; Uboweja, E.; Hays, M.; Zhang, F.; Chang, C.-L.; Yong, M. G.; Lee, J.; Chang, W.-T.; Hua, W.; Georg, M.; and Grundmann, M. 2019. MediaPipe: A Framework for Building Perception Pipelines. arXiv:1906.08172.
- Merdji, H.; Long, M. T.; Ostermann, M.; Herridge, M.; Myatra, S. N.; De Rosa, S.; Metaxa, V.; Kotfis, K.; Robba, C.; De Jong, A.; et al. 2023. Sex and gender differences in intensive care medicine. *Intensive care medicine*, 49(10): 1155–1167.
- Penoyer, D. A. 2010. Nurse staffing and patient outcomes in critical care: a concise review. *Critical care medicine*, 38(7): 1521–1528.
- Selvan, K.; Edriss, H.; Sigler, M.; and Tseng, J. 2014. Self-extubation in ICU patients. *The Southwest Respiratory and Critical Care Chronicles*, 2(8): 31–34.
- Sessler, C. N.; Gosnell, M. S.; Grap, M. J.; Brophy, G. M.; O’Neal, P. V.; Keane, K. A.; Tesoro, E. P.; and Elswick, R. 2002. The Richmond Agitation–Sedation Scale: validity and reliability in adult intensive care unit patients. *American journal of respiratory and critical care medicine*, 166(10): 1338–1344.
- Taylor, M. A.; Reynolds, C. M.; and Jones, R. 2021. Challenges and potential solutions for patient safety in an infectious-agent-isolation environment: a study of 484 COVID-19-related event reports across 94 hospitals. *Patient Safety*, 3(2): 45–62.
- Zhang, J.; Ma, G.; Peng, S.; Hou, J.; Xu, R.; Luo, L.; Hu, J.; Yao, N.; Wang, J.; and Huang, X. 2023. Risk factors and predictive models for peripherally inserted central catheter unplanned extubation in patients with cancer: prospective, machine learning study. *Journal of medical Internet research*, 25: e49016.