

Suicide Risk Assessment from AI-powered Video Surveillance: An Interpretable Framework for Prevention in Metro Stations

Safwen Naimi¹, Wassim Bouachir¹, Guillaume-Alexandre Bilodeau² and Brian Mishara³

¹Université TÉLUQ, QC, Canada

²Polytechnique Montréal, QC, Canada

³Université du Québec à Montréal, QC, Canada

{safwen.naimi, wassim.bouachir}@teluq.ca, guillaume-alexandre.bilodeau@polymtl.ca,
mishara.brian@uqam.ca

Abstract

Understanding and monitoring human behavior in metro stations play an important role in supporting suicide prevention efforts, where early identification of high-risk situations can enable timely intervention. This requires assessing suicide risk from a surveillance video by jointly reasoning about the behavior of each passenger, his/her spatial context, and temporal dynamics. However, this assessment using videos captured by surveillance cameras is challenging, as it demands accurate perception of human motion, understanding of platform geometry, and aggregation of heterogeneous behavioral cues over time. In this work, we formalize the task of Suicide Risk Assessment (SRA) in metro stations and introduce the first interpretable framework that addresses this challenge. Unlike approaches that focus on isolated subtasks or attempt to infer intent directly, our formulation assesses suicide risk from accumulated evidence by incorporating person tracking, activity recognition, semantic segmentation of the platform, and trajectory-driven risk heatmap modeling. By formalizing SRA as a distinct task and benchmarking a complete operational pipeline achieving 83.2% ROC-AUC on real surveillance data, this work highlights the complexity of suicide risk assessment and opens new directions for research on interpretable AI systems for social good.

1 Introduction

Suicide remains a major public health concern worldwide [Forte *et al.*, 2021; Rassy *et al.*, 2021]. The deployment of surveillance cameras at known suicide hotspots and high-risk public spaces, such as metro stations, has created opportunities for monitoring and early intervention [Anjali and Kumar, 2024; Bouachir *et al.*, 2018]. Existing vision-based systems, however, largely depend either on continuous human observation or on simple visual cues such as detecting the crossing of safety barriers to trigger emergency responses [Laliberte *et al.*, 2021]. These strategies often detect incidents at a late stage, limiting their preventive impact, while remain-

ing time-consuming and prone to human fatigue. Supporting human operators with automated systems that can highlight potentially high-risk situations in real time while remaining interpretable and ethically grounded has therefore become an important challenge at the intersection of artificial intelligence and social good [Balamurugan *et al.*, 2025; Gursel *et al.*, 2024].

Video surveillance infrastructures are already widely deployed in metro stations, offering a rich source of visual information for understanding passenger behavior and movement patterns. Several prior studies have investigated the analysis of videos to better understand behaviors preceding suicide attempts in metro and railway environments [Mishara *et al.*, 2016; Parsapoor *et al.*, 2023; Rådbo and Andersson, 2012; Silla and Luoma, 2012; Mackenzie *et al.*, 2018]. These works, often grounded in interdisciplinary collaborations between psychology, sociology, and transportation authorities, have identified recurring behavioral and spatial patterns, such as repeated pacing near the far-end zone of the platform, prolonged standing or walking on the yellow platform demarcation line, and looking toward the tunnel. Comparative investigations between attempters and non-attempters have further shown that risk-related behaviors are not defined by a single distinctive action, but by patterns of accumulation and repetition over time [Mishara *et al.*, 2016]. We believe that developing reliable systems for suicide risk assessment in metro stations requires an interdisciplinary approach, and that a key step forward is the integration of knowledge from researchers in suicide prevention with artificial intelligence advancement for effective real-time video analysis.

In this work, we formalize Suicide Risk Assessment (SRA) as a distinct computational task. Rather than predicting intent or labeling individuals, SRA aims to assess the level of suicide risk associated with observable behavior and spatial context over time for each passenger. This formulation aligns with interdisciplinary findings that view suicide risk as an emergent phenomenon, shaped by the interaction between individuals and their environment [Hightower and Grant, 2024; Gariépy *et al.*, 2024]. We propose an interpretable framework that integrates multiple complementary modules capable of producing continuous, individual-level suicide risk assessment grounded in spatial exposure, temporal persistence, and behavioral patterns, while avoiding reliance on appearance cues or personal attributes.

Interdisciplinary Collaboration. This work is the result of an interdisciplinary collaboration between computer scientists (at universities), researchers in the field of suicidology, and The Société de transport de Montréal (STM).

Contributions. We summarize our contributions as follows: (i) We introduce **SRA-Framework**, the first interpretable SRA pipeline for suicide risk assessment in metro stations, (ii) We design a set of spatial, temporal, and behavioral risk indicators aligned with findings in suicide prevention and show how their aggregation over time enables effective and interpretable suicide risk assessment, and (iii) We evaluate the proposed framework on ethically approved, real-world surveillance video data from metro stations, demonstrating meaningful discrimination between at-risk and control scenarios and providing qualitative and quantitative interpretability analyses.

2 Related Work

2.1 Behavioral Studies for Suicide Prevention in Metro Stations

The study of behaviors preceding a suicide attempt in public transportation stations, particularly through surveillance cameras, is a crucial research area for proactive suicide prevention. Several studies identified the behavioral signs that precede such acts in order to enable rapid intervention [Mishara *et al.*, 2016; Rådbo and Andersson, 2012; Silla and Luoma, 2012; Mackenzie *et al.*, 2018]. These studies report that attempters spend significantly more time on the yellow line and the far-end zone of the platform, compared to control passengers. In addition to spatial proximity, repeated movement patterns, such as back-and-forth pacing between the wall of the platform and the yellow line, as well as repetitive looking at the tunnel, have been shown to occur more frequently among attempters. Temporal factors play an important role because at-risk situations are characterized by the accumulation and repetition of risk-related behaviors over time. According to Mishara *et al.* [2016], 83% of attempters exhibited easily observable behaviours potentially indicative of an impending attempt, and 61% had two or more of these behaviours, while 75% had at least one of these behaviours. According to the same study, two behaviours: leaving an object on the platform and going back and forth between the wall and the yellow line, could identify 24% of attempters with no false positives. The other target behaviours could also be present in non-attempters, but with fewer occurrences. However, having two or more of these behaviours indicated a high likelihood of being at risk of attempting suicide.

2.2 Deep learning-based abnormal behavior recognition in Metro Stations

Deep learning approaches have been widely explored for abnormal behavior recognition from video surveillance, aiming to automatically identify unusual or potentially dangerous events in metro stations. These methods typically rely on spatio-temporal representations learned from raw video frames, optical flow, trajectories, or pose information, and are often formulated as anomaly detection or action classification problems. Naimi *et al.* [2025] introduced SSTAR, a

skeleton-based action recognition model designed to recognize pre-suicidal passenger activities. While SSTAR demonstrates strong performance in identifying semantically meaningful actions under challenging surveillance conditions, it does not account for spatial context and risk assessment. Zuo *et al.* [2025] formulates passenger abnormal behavior recognition as a video anomaly detection task, using a Swin Transformer-based generative model to reconstruct surveillance frames and detect abnormalities via reconstruction error, with additional motion refinement through optical flow. This approach achieves good performance but focuses on frame-level anomaly detection rather than interpretable, context-aware risk assessment. Chang *et al.* [2025] proposes a smart railway safety monitoring system that integrates deep learning-based segmentation and object detection to identify intrusions on railway tracks. Platform boundaries are automatically extracted using Mask R-CNN [He *et al.*, 2017], while YOLOv3 [Redmon and Farhadi, 2018] is employed to detect human and object intrusions under diverse environmental conditions. This approach focuses on intrusion detection and alerting rather than behavioral analysis or suicide risk assessment. Different from previous works, which largely focus on anomaly or intrusion detection, our SRA-Framework formulates suicide prevention as a long-term, context-aware risk assessment problem. We integrate multiple vision components: tracking, platform segmentation, activity recognition, and spatial homography to extract interpretable behavioral and spatial indicators that are accumulated over time. This enables continuous suicide risk monitoring under real surveillance conditions rather than instantaneous event detection.

3 Suicide Risk Assessment Task

We address the problem of suicide risk assessment in metro stations using continuous surveillance video streams. The objective is to assist preventive monitoring by assessing, for each observed passenger, a time-varying risk score that reflects their level of suicide risk based on observable behavior and spatial context.

Let $V = \{F_t\}_{t=1}^T$ denote a sequence of video frames acquired from a camera. At each time step t , the system detects and tracks a set of individuals $\mathcal{P}_t = \{p_i\}_{i=1}^{M_t}$. For each individual $p_i \in \mathcal{P}_t$, the task is to output a scalar suicide risk score $R_{p_i}(t) \in [0, 1]$, where higher values indicate elevated suicide risk. Formally, suicide risk is modeled as a function of accumulated behavioral and spatial indicators extracted over a temporal window $W_t = [t - \tau + 1, t]$ such as:

$$R_{p_i}(W_t) = \sigma(f(\mathbf{x}_{p_i}(W_t))), \quad (1)$$

where $\mathbf{x}_{p_i}(W_t) \in \mathbb{R}^d$ denotes a feature vector summarizing d observable behaviors and interactions between the individual and the platform environment over time, $f(\cdot)$ represents a supervised learning model, and $\sigma(\cdot)$ maps the output to the interval $[0, 1]$. Note that $R_{p_i}(W_t)$ does not represent suicidal intent, a psychological state, or a clinical diagnosis. Instead, it quantifies suicide risk, defined as the degree to which an individual’s observable behavior and spatial interaction with




Figure 1: **Architecture overview of our proposed SRA-Framework.** It takes a video as input and outputs the complete platform state and a suicide risk assessment for each person.

different areas of the platform may warrant preventive attention. It is influenced not only by what actions are performed, but also by where they occur and how long they persist. This formulation is consistent with operational safety practices, where decisions are informed by the accumulation of behavioral and contextual cues rather than isolated events.

4 Proposed SRA-Framework

Given a continuous surveillance video stream, our framework outputs a platform state representing a time-varying suicide risk score $R_{p_i}(W_t) \in [0, 1]$ for each tracked individual p_i . The overall architecture of our SRA-Framework is depicted in Fig. 1. First, human motion representations of each person on the platform are extracted using sequential detection, tracking, and pose estimation modules. Second, a skeleton-based activity recognition module uses pose dynamics to predict action labels of each person over short temporal windows. Third, a platform segmentation module captures the spatial structure of the station by segmenting the platform into relevant zones, which serve as a basis for identifying back-and-forth and repetitive movement behaviors. Fourth, a position risk heatmap generation module aggregates the platform-referenced heatmaps of multiple at-risk individuals over time, producing a position risk heatmap that reflects the most frequently traversed or occupied areas by persons at risk within the platform. Finally, a risk inference module aggregates per-person indicators into a feature vector $\mathbf{x}_{p_i}(t)$ and computes the suicide risk score $R_{p_i}(W_t)$. To simplify notation, in the following, we will denote an individual p_i as p .

4.1 Human Detection, Tracking, and Pose Estimation

We employ a pre-trained YOLOx [Jocher *et al.*, 2023] model as our human detector, without fine-tuning it on our dataset, since it already provides decent performance on human videos. We filter the model output to retain only the "person" class detections.

To leverage existing strong multi-object trackers, our SRA-Framework performs tracking in the image space based on bounding boxes. We employ ByteTrack [Zhang *et al.*, 2021] as our multi-object tracker, for its good performance and its ability to leverage both spatio-temporal and appearance cues. For each tracked individual, a pretrained HRNet [Wang *et al.*, 2019] model is applied to extract a set of 2D body keypoints within the corresponding bounding box. Pose estimates are accumulated over temporal windows to form pose sequences that capture fine-grained motion dynamics while remaining invariant to appearance and scale. This representation enables the system to reason about human movement patterns beyond coarse trajectory information.

As illustrated in Fig. 1, tracked positions are subsequently projected into a platform-referenced coordinate system using a homography matrix M , which enables the construction of trajectory-based heatmaps that encode where and how long each individual moves or remains within the platform.

4.2 Human Activity Recognition

Given the tracked pose sequences extracted in section 4.1, this module infers human activity patterns over temporal windows




Figure 2: **Overview of the Platform Semantics Modeling Process:** (a–b) Boundary and Offset Lines Estimation, and (c) Final Zone Partitioning.

W_t , providing high-level action cues for subsequent risk assessment. In our SRA-Framework, we employ SSTAR pretrained on the ARMM dataset [Naimi *et al.*, 2025], a skeleton-based spatio-temporal action recognition model specifically designed for surveillance scenarios. In our case, SSTAR is used to detect 3 passenger action labels: ‘LookTunnel’, ‘Walk’, and ‘Stand’. The output of this module consists of per-person activity labels associated with each temporal window. Importantly, the contribution of recognized activities in our system depends on their spatial context, duration, and repetition.

4.3 Platform Semantics and Zones Construction

To incorporate environmental context, we model platform semantics by partitioning the scene into spatial zones that support interpretable risk indicators, such as back-and-forth movements between the wall and the yellow line and explicit far-end zone entry events. Platform segmentation is obtained using YOLOv8n [Jocher *et al.*, 2023] in its semantic segmentation configuration, chosen for its favorable accuracy–efficiency trade-off. The model is trained offline on manually annotated images with pixel-level labels of the platform surface.

For zones construction, we estimate the two lateral platform boundaries corresponding to the wall-facing side (ℓ_L^*) and the track-facing side (ℓ_R^*) of the platform using the segmentation result. As shown in Fig. 2a, let T_L and T_R represent the top left point and the top right point of the platform contour respectively. Given the top platform edge segment $\ell_o = \overline{T_L T_R}$, we generate an inward parallel line segment ℓ_d at offset distance d using the unit normal vector. To ensure geometric validity within the platform extent, ℓ_d is clipped by intersecting it with the estimated left and right platform boundary line segments ℓ_L^* and ℓ_R^* , yielding the bounded segment $s_d = \overline{Q_L Q_R}$ where Q_L and Q_R denote the intersection points of ℓ_d with ℓ_L^* and ℓ_R^* respectively. These calibrations were established in consultation with experts from the transit authority.

Using the four platform corner points (T_L, T_R, B_L, B_R), two internal boundaries are defined by selecting points at fixed fractions of the top and bottom edges ($\alpha = 1/5$, $\beta = 1/7$) and connecting them, resulting in the line segments, ℓ_{left} and ℓ_{right} (Fig. 2b). The resulting boundaries

$\ell_{left} = \overline{T_L^{(\alpha)} B_L^{(\beta)}}$ and $\ell_{right} = \overline{T_R^{(\alpha)} B_R^{(\beta)}}$ will define the boundaries of wall-proximal and yellow-line-proximal zones. The final three zones are defined as shown in Fig. 2c:

- **Zone (A): Wall-Proximal Zone.** Corresponds to the area near the platform wall, farthest from the track. It is defined by the polygon $(T_L, T_L^{(\alpha)}, B_L^{(\beta)}, B_L)$.
- **Zone (B): Yellow-Line Proximal Zone.** Corresponds to a transitional space where passengers may approach the yellow line while remaining within designated safety boundaries. It is defined by the polygon $(T_R^{(\alpha)}, T_R, B_R, B_R^{(\beta)})$.
- **Zone (C): Platform Far-End Zone.** Corresponds to the area immediately adjacent to the platform edge, it is the nearest area to the tunnel defined by the polygon (T_L, T_R, Q_R, Q_L) .

4.4 Position Risk Heatmap Generation

We employ a homography transformation that maps image-plane coordinates to a platform-referenced coordinate system using the homography matrix M . This transformation allows tracked person positions to be expressed in a common spatial frame aligned with the actual platform geometry. As a result, spatial proximity, persistence, and movement patterns can be interpreted consistently across time and individuals.

Once projected into the platform-referenced space, each tracked individual is represented by a sequence of two-dimensional positions over time. The heatmap of each tracked individual is generated mathematically through a density accumulation and Gaussian smoothing process over its 2D trajectory points as shown in Fig. 3a. We define a heat density map $\mathcal{H}(x, y)$ as:

$$\mathcal{H}(x, y) = \sum_{j=1}^N \delta(x - x_j, y - y_j) \quad (2)$$

where (x_j, y_j) are the recorded 2D trajectory points of a person, N denotes the number of recorded trajectory points for a given individual, and $\delta(\cdot)$ is the Dirac delta representing one unit of heat per position.

Individual heatmaps provide information about the spatial behavior of each person, but suicide risk in metro platforms is




Figure 3: Illustration of Individual At-Risk Heatmaps and their aggregation into a Position Risk Heatmap.

also shaped by persistent environmental factors and recurrent patterns. A global representation is therefore required to capture shared areas that emerge across different at-risk individuals. To obtain this position risk heatmap, we aggregated K at-risk individual trajectory-based heatmaps from our dataset. This aggregation yields a scene-level spatial representation reflecting the collective distribution of different at-risk persons on the platform. The position risk heatmap $\bar{\mathcal{H}}$ is computed as:

$$\bar{\mathcal{H}}(x, y) = \frac{1}{K} \sum_{k=1}^K \mathcal{H}_k(x, y) \quad (3)$$

Regions that are repeatedly occupied or associated with prolonged dwell time naturally emerge in the heatmap as high-intensity areas, while transient regions remain low as depicted in Fig. 3b.

The aggregated position risk heatmap $\bar{\mathcal{H}}(x, y)$ is normalized in the range $[0, 1]$ to obtain a position risk score PR_p . This position risk score is later integrated to support the final suicide risk assessment stage.

4.5 Suicide Risk Assessment

Based on the behavioral studies for suicide prevention in metro stations discussed in Section 2.1, we derive eight interpretable suicide risk indicators for each person p from the outputs of the previous sections:

1. **Position risk score (PR_p)**. As described in Section 4.4, it is computed by averaging the position risk heatmap $\bar{\mathcal{H}}$ values along an individual’s trajectory, capturing whether a person occupies spatial locations associated with recurrent exposure and quantifying this occupation.
2. **Walk/Stand on the yellow line (Cr_p)**. This binary indicator denotes whether an individual is detected performing a walking or standing action while positioned on the yellow line. It is activated when a person’s left or right leg keypoint crosses the right boundary of Zone B and is simultaneously classified as walking or standing by the activity recognition module (Section 4.2).
3. **Yellow line crossings (NCr_p)**. It counts the number of distinct events in which the right or left leg keypoints of an individual cross the right boundary of Zone B.
4. **Time on the yellow line (TY_p)**. It accumulates the total duration spent on the yellow line region in seconds.

5. **Longest yellow line segment (LY_p)**. It records the maximum continuous duration spent on the yellow line region without interruption.
6. **Back-and-forth between wall and yellow line count (BF_p)**. It measures repeated transitions between Zone A and Zone B derived from Section 4.3. A transition is recorded when the right and left leg keypoints of an individual cross one of the two zones, move to the other zone, and subsequently return to the initial one: ‘A $\rightarrow$ B $\rightarrow$ A’ or ‘B $\rightarrow$ A $\rightarrow$ B’ transitions.
7. **Look Tunnel (LT_p)**. This binary indicator is obtained from the activity recognition module (Section 4.2) and indicates whether the individual is detected repeatedly orienting his upper body toward the tunnel area.
8. **Enter the far-end zone of the platform (E_p)**. This binary indicator determines whether the individual has entered the platform far-end zone (Zone C).

Given the set of risk indicators defined above, we assess suicide risk using an XGBoost supervised model [Chen and Guestrin, 2016] operating on accumulated features over a temporal window. We adopt XGBoost for its robustness under limited and imbalanced data conditions, making it well-suited for our individual-level suicide risk assessment [Veldar *et al.*, 2023]. As defined in Eq. (1), for an individual p , a feature vector $\mathbf{x}_p(W_t) \in \mathbb{R}^d$ is extracted over W_t as:

$$\mathbf{x}_p(W_t) = [PR_p(W_t), Cr_p(W_t), NCr_p(W_t), TY_p(W_t), LY_p(W_t), BF_p(W_t), LT_p(W_t), E_p(W_t)] \quad (4)$$

where each component corresponds to one of the proposed suicide risk indicators accumulated within the window W_t . This feature vector is then provided as input to our XGBoost, which outputs a continuous suicide risk score $R_p(W_t)$ between 0 and 1.

5 Experimental Evaluation

5.1 Dataset Description

With approval from the Société de transport de Montréal (STM) and the Research ethic Institutional Review Board of our universities, we accessed surveillance video recordings of suicide attempts in the metro stations of Montréal. They provided 66 five-minute surveillance video recordings of cases before an attempt, and 56 recordings from the same cameras at the same time of day, and day of week, but when no suicide attempt occurred. For the suicide risk assessment task, we created a dataset with annotations provided at the individual level, including their extracted risk indicators, with the assistance of psychology experts and researchers in the field of suicidology. Individuals are labeled according to whether they originate from an at-risk or control scenario, without frame-level intent annotations. Using this procedure, we obtained a total of 256 individual instances, comprising 190 control subjects and 66 at-risk subjects. Data splits are performed at the video level, ensuring that all individuals extracted from the same recording belong to the same subset. The dataset was split into 75% training and 25% test, while preserving class proportions between at-risk and control.




Figure 4: Feature importance of the eight risk indicators used for suicide risk assessment, computed from our trained XGBoost model. Higher F-scores indicate more frequent usage of a feature.

5.2 Implementation Details

Our XGBoost model is trained with 300 boosting iterations and a learning rate of 0.05. Tree depth is limited to four to control model complexity, while subsampling of training instances and features is applied with ratios of 0.85 to improve generalization. Class imbalance between control and at-risk individuals is addressed by weighting the positive class proportionally to its inverse frequency in the training set. The model is optimized using the binary logistic objective and primarily evaluated using the ROC-AUC metric, which measures the ability of the model to discriminate between at-risk and control individuals independently of any decision threshold. In addition, we report Sensitivity, which reflects correct identification of at-risk cases, and Specificity which reflects correct identification of control cases. We also report the false positive rate (FPR), corresponding to false alarms, and the false negative rate (FNR), corresponding to missed detection rate. These metrics are computed by applying a fixed decision threshold of $R_p > 0.2$, chosen to favor early detection of at-risk situations in line with suicide prevention objectives.

5.3 Quantitative Results

Table 1 reports 10-Fold Cross-Validation results of the proposed suicide risk assessment framework under different controlled configurations, where detection/tracking and activity recognition modules are alternatively replaced by ground-truth annotations to assess the robustness of our framework and to quantify the impact of different modules on the final decision. The results show that replacing the predicted detection/tracking module in the fully automatic configuration with ground-truth annotations increases ROC-AUC from 0.832 ± 0.095 to 0.919 ± 0.054 , while substantially improving sensitivity ($0.631 \rightarrow 0.800$) and reducing false negative rates ($0.369 \rightarrow 0.200$). In contrast, substituting the activity recognition module with ground-truth labels yields more moderate gains, with ROC-AUC increasing from 0.832 ± 0.095 to 0.893 ± 0.051 . The remaining gap between fully automatic configuration (0.832 ± 0.095) and the upper bound configuration (0.924 ± 0.055) highlights room for future improvements to the indicator computation modules. As well, it highlights the need for researchers in suicidology to continue studying possible visual indicators that can be actionable with com-




Figure 5: Model interpretability analysis using SHAP values, highlighting how each indicator influences suicide risk predictions.

puter vision and machine learning algorithms.

Beyond predictive performance, we analyze model interpretability to better understand the contribution of each risk indicator. Fig. 4 reports the feature importance scores obtained from the trained XGBoost model for the fully automatic configuration, which reflects how frequently a feature is used to split the data across the ensemble of trees. The results show that indicators related to direct interaction with the yellow line play a dominant role in suicide risk assessment. In particular, crossing the yellow line emerges as the most influential feature (≈ 8.1 F-Score), followed by the number of yellow line crossings and the back-and-forth movement count. The time spent on the yellow line and the position risk score also contribute substantially, highlighting the importance of prolonged presence in high-risk regions rather than isolated events. Other behavioral cues, such as entering the far-end zone of the platform and looking at the tunnel, provide complementary information but are less decisive when considered independently. These findings are consistent with prior psychological and behavioral studies.

Fig. 5 presents the SHAP summary plot, illustrating both the relative importance of each risk indicator and the directional influence of their values on the predicted suicide risk score. Consistent with the feature importance analysis, position risk score and direct interaction with the yellow line emerge as dominant contributors. Higher values of the position risk score are associated with positive SHAP values, indicating a strong increase in predicted risk when individuals occupy spatial regions frequently associated with prior at-risk trajectories. Similarly, crossing the yellow line shows a clear separation, where positive instances (high feature values) consistently push the suicide risk score upward. Indicators capturing repetitive and persistent behaviors, such as back-and-forth movements, number of yellow line crossings, and time spent on the yellow line, also exhibit predominantly positive SHAP contributions when their values increase.

5.4 Qualitative Results

Fig. 6 illustrates two platform states predicted by our SRA-Framework. The first video illustrated in Fig. 6a presents three persons on the platform, our system correctly assigns a high risk score ($R_0=0.98$), driven by a combination of a prolonged crossing of the yellow line (11 seconds) and elevated position risk score along the position risk heatmap. Our sys-

Configuration	Modules		Statistical Measures ($\pm$ SD)				
	Det./Track.	Act Rec	ROC-AUC	Sensitivity	Specificity	False Alarm Rate (FPR)	Missed Det. Rate (FNR)
GT-Assisted	YOLOx/Bytetrack	GT	0.893 ± 0.051	0.631 ± 0.228	0.913 ± 0.056	0.087 ± 0.056	0.369 ± 0.228
GT-Assisted	GT	SSTAR	0.919 ± 0.054	0.800 ± 0.187	0.915 ± 0.046	0.085 ± 0.046	0.200 ± 0.187
Upper Bound	GT	GT	0.924 ± 0.055	0.867 ± 0.113	0.929 ± 0.064	0.071 ± 0.064	0.133 ± 0.113
Fully Automatic	YOLOx/Bytetrack	SSTAR	0.832 ± 0.095	0.631 ± 0.228	0.902 ± 0.071	0.098 ± 0.071	0.369 ± 0.228

Table 1: 10-Fold Cross-Validation results of our SRA-Framework under different module configurations. In the GT-Assisted configurations, we replaced Detection/Tracking (Det./Track.) or Activity recognition (Act Rec) modules with GT annotations. In the Upper Bound configuration, we exclusively used GT annotations for SRA. SD corresponds to the standard deviation.



Figure 6: **Qualitative Results.** Output predictions of two frames from surveillance video streams.

tem also successfully assigned lower suicide risk scores for control individuals ($R_2=0.12$ and $R_3=0.14$). In the second video illustrated in Fig. 6b, our system correctly assigns a high risk score of $R_0=0.92$ to the at-risk individual (ID 0), while correctly estimating substantially lower risk scores for control individuals ($R_1=0.38$ and $R_3=0.06$). This distinction arises from the accumulation of multiple risk indicators for ID 0, including prolonged crossing of the yellow line, a look-tunnel action, and sustained presence in spatial regions associated with elevated risk in the position risk heatmap.

6 Limitations and Future Work

Despite encouraging results, several limitations should be acknowledged. First, the dataset used in this study is limited in size. Although this reflects real-world deployment conditions, larger and more diverse dataset would be beneficial. Second, for the proposed framework, errors introduced at early stages, particularly in detection, tracking, and activity recognition, can propagate to downstream risk indicators. Our controlled evaluations using ground-truth substitutions in Table 1 show that further improvements in module choice could significantly enhance the overall performance. Third, the set of risk indicators considered in our SRA-Framework does not cover all behavioral cues identified in psychological and behavioral studies, such as leaving an object on the platform. Waiting on the platform until one or more trains passed was not covered in our study because of using only 5-min videos before the train arrived.

Future research directions include improving tracking under crowded conditions and incorporating additional interpretable behavioral indicators. Ultimately, integrating human-in-the-loop feedback and evaluating the framework in real-time operational settings with the STM represent important steps toward practical deployment.

7 Conclusion

Our work introduces the first complete framework for Suicide Risk Assessment (SRA) in metro stations, integrating perception, behavioral analysis, spatial reasoning, and interpretable risk inference within a unified pipeline. Unlike prior works that address isolated components such as detection, tracking, or activity recognition, our approach benchmarks a full operational pipeline whose outputs directly support safety-oriented decision making in real-world surveillance environments. Beyond performance, our study emphasizes interpretability, showing that risk assessments are driven by intuitive indicators aligned with established behavioral and spatial risk factors. This positions the proposed framework as a meaningful bridge between AI-based surveillance systems and interdisciplinary research on suicide prevention.

Acknowledgements

This work was supported by the Canada Research Chair in Artificial Intelligence for Suicide Prevention (CRC-2023-00036).

Ethical Statement

This study was conducted in accordance with applicable ethical guidelines and received approval from the relevant institutional ethics committee. All video data were handled under confidentiality protocols and anonymized through systematic blurring of identifiable visual information prior to analysis. The data were used exclusively for suicide prevention research, securely stored, and accessible only to authorized researchers.

The literature documents differences in suicidal behaviors across sex and age groups, as well as algorithmic biases in computer vision systems related to individuals' characteristics. However, the size of our dataset and the quality of the recordings did not allow for a differentiated analysis by sex and age to assess potential biases. In future work, access to a larger dataset collected with a new generation of cameras will enable disaggregated analyses by sex and age groups to examine behavioral variations.

References

- [Anjali and Kumar, 2024] T Anjali and B. Rushi Kumar. Spatial analysis of multivariate factors influencing suicide hotspots in urban tamil nadu. *Journal of Affective Disorders Reports*, 2024.
- [Balamurugan *et al.*, 2025] Merlin Balamurugan, Kumar Shanmugasamy, and Saranya Balaguru. Humans in the loop, lives on the line: Ai in high-risk decision making. *European Journal of Computer Science and Information Technology*, 2025.
- [Bouachir *et al.*, 2018] Wassim Bouachir, Rafik Gouiaa, Bo Li, and Rita Noumeir. Intelligent video surveillance for real-time detection of suicide attempts. *Pattern Recognit. Lett.*, 110:1–7, 2018.
- [Chang *et al.*, 2025] Chin-Chieh Chang, Kai-Hsiang Huang, Tsz-Kin Lau, Chung-Fah Huang, and Chun-Hsiung Wang. Using deep learning model integration to build a smart railway traffic safety monitoring system. *Scientific Reports*, 15, 2025.
- [Chen and Guestrin, 2016] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016.
- [Forte *et al.*, 2021] Alberto Forte, Giuseppe Sarli, Lorenzo Polidori, David Lester, and Maurizio Pompili. The role of new technologies to prevent suicide in adolescence: A systematic review of the literature. *Medicina*, 57, 2021.
- [Gariépy *et al.*, 2024] Geneviève Gariépy, Rifat Zahan, Nathaniel D. Osgood, Benjamin Yeoh, Eva Graham, and Heather Orpana. Dynamic simulation models of suicide and suicide-related behaviors: Systematic review. *JMIR Public Health and Surveillance*, 10, 2024.
- [Gursel *et al.*, 2024] Ezgi Gursel, Mahboubah Madadi, Jamie Baalis Coble, Vivek Agarwal, Vaibhav Yadav, Ronald L. Boring, and Anahita Khojandi. The role of ai in detecting and mitigating human errors in safety-critical industries: A review. *Reliab. Eng. Syst. Saf.*, 256:110682, 2024.
- [He *et al.*, 2017] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask r-cnn. 2017.
- [Hightower and Grant, 2024] Heath H. Hightower and Morgan J. Grant. Proposing an individual-in-contexts model for reimagining suicide screening, assessment, and intervention in black communities. *Journal of Human Services*, 2024.
- [Jocher *et al.*, 2023] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Ultralytics yolov8 <https://github.com/ultralytics/ultralytics>. 2023.
- [Laliberte *et al.*, 2021] Avery Z. Laliberte, Brandon Roth, Beau Edwards, and Jason I. Chen. Suicide screening and risk assessment in the emergency department: Case review of a suicide attempt survivor. *Journal of emergency nursing*, 2021.
- [Mackenzie *et al.*, 2018] Jay-Marie Mackenzie, Jo Borrill, Emily Hawkins, Bob Fields, Ian Kruger, Ian Noonan, and Lisa Marzano. Behaviours preceding suicides at railway and underground locations: a multimethodological qualitative approach. *BMJ Open*, 8, 2018.
- [Mishara *et al.*, 2016] Brian L. Mishara, Cécile Bardon, and Serge Dupont. Can cctv identify people in public transit stations who are at risk of attempting suicide? an analysis of cctv video recordings of attempters and a comparative investigation. *BMC Public Health*, 16, 2016.
- [Naimi *et al.*, 2025] Safwen Naimi, Wassim Bouachir, Guillaume-Alexandre Bilodeau, and Brian L. Mishara. Sstar: Skeleton-based spatio-temporal action recognition for intelligent video surveillance and suicide prevention in metro stations. *2025 IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACVW)*, pages 765–775, 2025.
- [Parsapoor *et al.*, 2023] Mahboobeh Parsapoor, Jacob W. Koudys, and Anthony C. Ruocco. Suicide risk detection using artificial intelligence: the promise of creating a benchmark dataset for research on the detection of suicide risk. *Frontiers in Psychiatry*, 14, 2023.
- [Rådbo and Andersson, 2012] Helena Rådbo and Ragnar Andersson. Patterns of suicide and other trespassing fatalities on state-owned railways in greater stockholm; implications for prevention. *International Journal of Environmental Research and Public Health*, 9:772 – 780, 2012.
- [Rassy *et al.*, 2021] Jessica Rassy, Cécile Bardon, Luc Dargis, Louis-Philippe Côté, Laurent Corthésy-Blondin, Carl-Maria Mörch, and Réal Labelle. Information and communication technology use in suicide prevention: Scoping review. *Journal of Medical Internet Research*, 23, 2021.
- [Redmon and Farhadi, 2018] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *ArXiv*, abs/1804.02767, 2018.
- [Silla and Luoma, 2012] Anne Silla and Juha Luoma. Main characteristics of train-pedestrian fatalities on finnish railroads. *Accident; analysis and prevention*, 45:61–6, 2012.

- [Velarde *et al.*, 2023] Gissel Velarde, Anindya Sudhir, Sanjay Deshmane, Anuj Deshmunkh, Khushboo Sharma, and Vaibhav Joshi. Evaluating xgboost for balanced and imbalanced data: Application to fraud detection, 2023.
- [Wang *et al.*, 2019] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43:3349–3364, 2019.
- [Zhang *et al.*, 2021] Yifu Zhang, Pei Sun, Yi Jiang, Dongdong Yu, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. *ArXiv*, abs/2110.06864, 2021.
- [Zuo *et al.*, 2025] Jing Zuo, Zexin Li, Ming He, and Qiaoli Yang. The research on recognition of passenger abnormal behavior in urban rail transit stations based on conditional generative adversarial networks. *Alexandria Engineering Journal*, 2025.