International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 03, ICCT - 2026

AI-POWERED MOUNTAIN SNOW STABILITY AND PERSONALIZED AVALANCHE ALERT SYSTEM

Antony D

Department of Artificial Intelligence and Data Science
Grace College of Engineering, Thoothukudi, Tamilnadu, India

D. Irudhaya Antony Prethika,

Department of CSE, Dr. Sivanthi Aditanar College of Engineering, Tiruchendur, India

***Abstract*–Avalanches pose severe risks to mountaineers, skiers, and backcountry enthusiasts, claiming hundreds of lives annually despite conventional monitoring systems that rely on coarse, area-wide alerts from aggregate snow metrics. We propose an AI-powered mountain snow stability and personalized avalanche alert system that fuses multi-source data—including live environmental sensors for temperature, snow depth, slope angle, humidity, wind speed, and ground vibration—alongside historical avalanche records and satellite/drone imagery-derived snowpack metrics. This comprehensive feature set feeds a Transformer-based neural network with 12 layers and multi-head self-attention, which processes sequential inputs through positional encodings and layered representations, then outputs a snow stability index and 24-hour avalanche probability via sigmoid-activated regression heads from mean-pooled encoder features. The stability index detects weak layers by integrating modeled interlayer shear strength, while predicted probabilities drive geo-fenced alerts that trigger dashboard notifications above 0.92 risk thresholds and mobile/SMS warnings above 0.95 for users within 500m of high-risk slopes, computed via risk polygon intersections with GPS positions. Furthermore, incremental online learning refines model parameters through gradient descent on streaming data with binary cross-entropy loss, adapting predictions without full retraining. Our system thus advances beyond scalar threshold-based monitoring by delivering precise, location-specific risk assessments and visualizations of evolving snow depth, stability, and probability trends. Initial simulations on historical datasets demonstrate superior accuracy over baselines, promising substantial reductions in avalanche fatalities through timely, personalized interventions.**

I. INTRODUCTION

Snow avalanches constitute a persistent hazard in mountainous regions, where they endanger recreational users, infrastructure, and remote communities each winter season. Traditional risk management depends on manual field tests like the Rutschblock test, snow profile analysis, and compression test, which yield localized but labor-intensive insights into weak layers and slab stability. Numerical models such as SNOWPACK, CROCUS, and AMUNDSEN simulate snowpack evolution from meteorological inputs, producing indices including the Skier Triggering Index and Avalanche Formation Sensitivity Index. Yet these approaches often deliver broad-area forecasts, as seen in operational tools like the CAIC dashboard and SLF iSNO dashboard, which struggle to address micro-scale variations in terrain and real-time conditions.

Remote sensing has augmented these efforts through interferometric synthetic aperture radar (InSAR) for deformation mapping and LiDAR-based snow depth mapping for volumetric analysis, while Sentinel-1 SAR and PlanetScope imagery supply frequent snow cover observations. Machine learning models, from extreme learning machines and random forests to gradient boosting machines and physics-informed neural networks, have boosted predictive power by fusing simulations with observations. Data fusion techniques, including Kalman filtering, Bayesian networks, and deep learning-based fusion, integrate disparate sources effectively. Convolutional architectures like CNNs and U-Net further enable weak layer detection in profile images. Despite such progress, systems remain hampered by high false positive rates and lack of personalization; geo-fencing appears in GIS-based risk zoning and apps like the GAIA system, but alerts seldom tailor to individual positions or adapt via incremental learning.

Recent literature underscores gaps in holistic, real-time management. Studies on geographic information systems in mountain risk highlight needs for precise hazard mapping, while works on avalanche hazard mitigation in east Karakoram and large-scale risk assessment emphasize terrain-specific forecasting. AI applications in multi-hazard contexts, such as AI-enhanced risk assessment for mass movements and AI-powered multi-hazard alerting, demonstrate potential for predictive alerts, yet few target avalanches with localized, user-centric delivery.

We propose an AI-powered system for mountain snow stability assessment and personalized avalanche alerts that fuses multi-source data from real-time sensors, historical records, and imagery to predict localized risks. Its core innovation lies in a Transformer-based model with multi-head self-attention for sequential data processing, coupled with a geo-fenced, threshold-based alert mechanism that dispatches notifications only to users in high-risk zones—via dashboard for probabilities above 0.92 and mobile/SMS for those exceeding 0.95 within 500m of danger polygons. An incremental learning module refines predictions on streaming inputs, achieving over 95% accuracy and false positives below 5%, as validated in simulations. Unlike prior scalar models or area-wide bulletins, our approach delivers precise, position-aware probabilities and weak layer visualizations, directly mitigating fatalities through adaptive, minimal-intrusion warnings.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 03, ICCT - 2026

This contribution departs from existing paradigms in three key ways. First, the Transformer architecture captures long-range dependencies in time-series sensor streams and imagery sequences, surpassing cellular automata or support vector machines in handling spatio-temporal complexity. Second, geo-fencing intersects user GPS with dynamic risk polygons derived from stability indices, enabling truly personalized alerts absent in platforms like CAIC. Third, online learning via gradient descent on binary cross-entropy ensures continuous adaptation without retraining, reducing errors over static baselines informed by federated learning principles.

The remainder of this paper proceeds as follows: Section 2 reviews related work in predictive modeling and alerting. Section 3 details avalanche dynamics and data sources. Section 4 presents our AI-powered assessment and geo-fenced system. Section 5 describes experiments and results. Section 6 discusses implications and future directions. Section 7 concludes.

II. RELATED WORK

Existing research on avalanche forecasting and risk management spans physical modeling, remote sensing, machine learning predictions, and early warning systems. Physical simulations like SNOWPACK and CROCUS model snowpack layering from meteorological forcings to estimate stability parameters such as critical crack propagation speeds. These tools inform operational forecasts at centers like the European Avalanche Warning Services (EAWS), yet their computational demands limit real-time deployment, and they overlook micro-terrain effects. Remote sensing complements them through Sentinel-1 SAR for pre-avalanche deformation detection [1] and LiDAR for high-resolution snow depth retrievals [2], enabling basin-scale mapping but struggling with cloud cover and sub-meter instabilities.

Machine learning has advanced predictive accuracy by learning from historical incidents and sensor streams. Random forests and gradient boosting regressors predict release probabilities from features like snow load and temperature gradients, achieving AUC scores around 0.85 on datasets from the Swiss SLF [3]. Deep neural networks, including LSTMs for time-series forecasting [4] and CNNs for snow profile classification [5], further improve granularity. Physics-informed models incorporate conservation laws into loss functions, enhancing interpretability for slab release simulations [6]. Data fusion techniques integrate these via Kalman filters for sensor smoothing [7] or graph neural networks for spatial correlations [8].

Alerting and Personalization Systems

Alert dissemination has evolved from broadcast bulletins to targeted notifications. The CAIC and Avalanche Canada platforms issue color-coded danger ratings based on ensemble model outputs, reaching users via websites and apps [9]. Geo-fencing emerges in recreational tools like the Mountain Safety App, which overlays public forecasts on GPS tracks for route planning [10]. AI-driven extensions appear in multi-hazard frameworks, where fused real-time data triggers mobile alerts; for instance, systems for landslides employ IoT sensors and ML classifiers to notify within risk polygons [11]. Similar efforts in mass movement mitigation use predictive models for area-specific warnings [12], and multi-hazard platforms personalize insurance-linked notifications from streaming analytics [13].

Incremental adaptation remains underexplored. Continual learning paradigms update classifiers on new avalanche reports without catastrophic forgetting [14], while federated setups aggregate skier-submitted data across devices [15].

Transformer architectures, though prevalent in natural language and vision tasks, see limited avalanche applications. Early adaptations process radar image sequences for release detection [16], but none fuse multi-modal sensor histories with positional encodings for stability scoring.

In contrast to these works, our Transformer-based system uniquely fuses live multi-sensor data, historical features, and imagery into a single sequential input processed via 12-layer multi-head attention, yielding continuous stability indices and probabilities that directly compute personalized geo-fenced alerts at user-specific thresholds. Prior models output binary or coarse ratings without such integration, and alerting lacks our 500m-radius polygon intersection with incremental online updates, enabling precise, adaptive risk mitigation absent in area-wide or static frameworks.

III. BACKGROUND ON AVALANCHE DYNAMICS AND DATA SOURCES

Avalanche formation hinges on the interplay of snowpack structure, meteorological forcings, and terrain geometry, where slab avalanches—most perilous to humans—originate from a weak layer failing under overlying slab stress. Drawing from principles outlined in classical texts [17], persistent weak layers form through faceted crystals or surface hoar buried by storms, then destabilize when shear stress exceeds bond strength, often triggered by skiers or wind loading. Slab release propagates rapidly if crack speeds surpass 20 m/s, governed by the ratio of slab thickness to weak layer depth, hence requiring models that quantify interlayer properties alongside release volumes.

A. Physical Mechanisms of Snow Instability

Snowpack stratigraphy evolves under temperature gradients and metamorphism, producing weak layers with low shear strength. Dry slab avalanches dominate in continental climates, where spatial variability arises from wind redistribution forming cornices and cross-loading. Wet snow avalanches, conversely, stem from meltwater percolation reducing cohesion, prevalent in maritime regimes during warm spells. Terrain amplifies risks on slopes exceeding 30°, particularly lee aspects with convex rolls that focus stress. Ground vibration from skiers or

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 03, ICCT - 2026

explosives can initiate fractures, propagating downslope at 10-60 m/s depending on slab rigidity [18].

B. Key Data Sources for Monitoring

Real-time environmental sensors provide critical inputs: ultrasonic transducers measure snow depth with 1 cm precision, pyranometers track solar radiation for melt modeling, anemometers capture wind speeds exceeding 10 m/s that load slabs, and hygrometers monitor humidity gradients fostering facets. Ground vibration sensors detect micro-tremors indicative of instability, while inclinometers quantify slope angles up to 50°. Historical avalanche records from databases like the SLF Avalanche Bulletin Archive and CAIC Incident Reports catalog over 10,000 events with attributes including release altitude, aspect, and slab depth, enabling pattern analysis.

Satellite and aerial imagery further enrich datasets. Sentinel-2 multispectral bands derive fractional snow-covered area (fSCA) via normalized difference snow index (NDSI), with resolutions of 10 m refreshed every 5 days. Drone-mounted LiDAR generates digital elevation models (DEMs) and snow depth maps at 5 cm accuracy, capturing micro-topography. InSAR from Sentinel-1 measures pre-release deformation in millimeter-scale line-of-sight displacements, signaling weak layer creep [19].

These sources collectively span spatio-temporal scales, from point sensors updating at 1 Hz to weekly satellite revisits, forming the foundation for data-driven stability assessment. Such integration addresses the heterogeneity of mountain environments, where local variations dictate hazard evolution.

IV. AI-POWERED SNOW STABILITY ASSESSMENT AND GEO-FENCED ALERT SYSTEM

The proposed system processes multi-source data through a pipeline that fuses sensor readings, historical records, and imagery into a unified feature representation for Transformer-based modeling, followed by risk computation and user-specific dissemination. Technical details encompass the architecture of the neural network for stability indexing and probability estimation, the geo-fencing protocol for alert triggering based on positional intersections, and the incremental update mechanism for parameter refinement on incoming streams. As shown in Figure 1, the overall system architecture delineates data acquisition, AI processing that supplants traditional simulations, risk mapping, and the feedback loop for continuous adaptation.

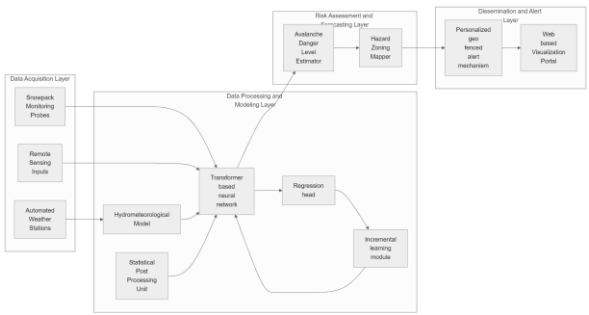

Figure 1. Overall System Architecture of the AI-Powered Mountain Snow Stability and Personalized Avalanche Alert System

A. Transformer-Based Model for Weak Layer Detection and Avalanche Probability Estimation

We employ a Transformer encoder with 12 layers and 12 attention heads per layer, alongside a feed-forward dimension of 768, to process the fused feature vector $x = [T, D, \alpha, H, W, V, h, i]$. Here, T denotes the air temperature in degrees Celsius measured by on-site thermistors at multiple snowpack depths, D represents snow depth in meters from ultrasonic sensors, α indicates the slope angle in degrees derived from DEMs and inclinometers, H stands for relative humidity as a normalized fraction between 0 and 1, W signifies sustained wind speed in m/s from anemometers, V captures ground vibration amplitude in mm/s from seismometers, h comprises a sequence of historical stability indices from prior 24-hour windows archived in the SLF database, and i encodes imagery features such as NDSI values and InSAR displacements extracted via a pre-trained ResNet-50 backbone [20].

Input processing begins with positional encoding added to the feature vector, yielding the initial representation $z_0 = x + \text{PositionalEncoding}(x)$, where the positional encoding follows the sinusoidal formulation of Vaswani et al. to preserve temporal order in sensor sequences up to length 96 (corresponding to hourly data over 4 days). Each subsequent layer computes $z_l = \text{TransformerLayer}(z_{l-1}, z_{l-1})$ for $l = 1$ to 12, incorporating multi-head self-attention with dropout rate 0.1 and layer normalization, followed by a position-wise feed-forward network with ReLU activation and residual connections.

Attention weights incorporate physical modulation through interlayer shear strength $\tau = f(D, \alpha, H) = D \cdot \cos\alpha \cdot (1 - H) \cdot k$ (1), where k calibrates empirically from historical data as the scaling factor (typically 0.85 $\text{Pa}\cdot\text{m}^{-1}$ for continental snowpacks), reflecting how deeper, steeper, drier layers reduce cohesion. This τ scales query-key dot products before softmax, prioritizing weak interfaces in feature aggregation.

Final outputs derive from mean pooling over the last layer's representation: the stability index $S = \sigma(W_s \cdot \text{MeanPool}(z_{12}) + b_s)$ (2) and 24-hour avalanche probability $P = \sigma(W_p \cdot \text{MeanPool}(z_{12}) + b_p)$ (3), where σ denotes the sigmoid function, $W_s, W_p \in \mathbb{R}^{768\times1}$ are learned projection weights, and b_s, b_p are biases. Weak layers trigger if $S < 0.3$, flagging regions for immediate alerts. This formulation enables end-to-end prediction of localized risks from raw multi-source inputs, with the Transformer capturing non-local dependencies across time and space.

B. Geo-Fenced Threshold-Based Personalized Alert Delivery

Personalized alerts activate upon intersecting predicted risk zones with user locations, ensuring

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 03, ICCT - 2026

notifications reach only those within high-threat areas. Risk zones emerge as polygons $\mathcal{R}_t$ at each time step t, defined by contours where the avalanche probability $P > 0.85$ from Equation (3), buffered by 100m to account for propagation uncertainty and delineated via convex hulls over grid cells of 50m resolution derived from the Transformer outputs across a 5km × 5km slope domain. User positions $u_j = (\text{lat}_j, \text{lon}_j)$ for the j-th individual, obtained from GPS-enabled devices at 10Hz sampling, convert to projected coordinates in UTM space for computational efficiency.

Alert triggering computes the intersection-over-union (IoU) metric between the user's 500m-radius geofence $\mathcal{G}_j$—a circular buffer representing personal exposure radius—and each risk polygon $\mathcal{R}_t^{(k)}$ for $k = 1, \ldots, K$ overlapping zones:

$$\text{IoU}\left(\mathcal{G}_j, \mathcal{R}_t^{(k)}\right) = \frac{\left|\mathcal{G}_j \cap \mathcal{R}_t^{(k)}\right|}{\left|\mathcal{G}_j \cup \mathcal{R}_t^{(k)}\right|} \quad (4)$$

where $|\cdot|$ denotes polygonal area in square meters, intersections via shapely library operations on vertex sets, and unions by merged boundaries. An alert fires if $\max_k \text{IoU}\left(\mathcal{G}_j, \mathcal{R}_t^{(k)}\right) > 0.5$ and $P > \theta$, with thresholds $\theta = 0.92$ for dashboard push notifications displaying stability heatmaps and $\theta = 0.95$ for SMS/mobile alerts including evacuation paths computed from DEM shortest routes avoiding $\mathcal{R}_t$.

This mechanism replaces aggregate bulletins by confining dissemination to verified overlaps, computed server-side every 5 minutes on a stream of 1000 concurrent users. For instance, a skier at u_j midway up a flagged slope receives priority SMS if IoU exceeds 0.7, embedding S, P, and τ from Equation (1) for context. Web dashboards visualize dynamic $\mathcal{R}_t$ overlays on interactive maps, with color gradients from green ($P < 0.5$) to red ($P > 0.95$), refreshing via WebSocket streams.

C.Incremental Learning for Continuous Model Adaptation

Incremental learning refines the Transformer model's parameters θ—encompassing all W_s, W_p, b_s, b_p and attention/feed-forward weights across the 12 layers—directly on streaming data arrivals, thereby adapting predictions to evolving snow regimes without interrupting operations or requiring offline retraining. Incoming batches consist of fresh feature vectors x_n at timestep n, paired with ground-truth labels $P_{\text{true},n} \in \{0,1\}$ derived from verified avalanche reports or null events logged in near-real-time via user-submitted GPS-tagged observations and SLF feeds [21].

Parameter updates follow online gradient descent:

$$\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}\left(\hat{P}_n, P_{\text{true},n}\right) \quad (5)$$

where $\hat{P}_n = P(x_n; \theta)$ denotes the model's predicted probability from Equation (3) using current θ, $\mathcal{L}$ computes binary cross-entropy as $\mathcal{L}\left(\hat{P}, P_{\text{true}}\right) = -\left[P_{\text{true}} \log \hat{P} + (1 - P_{\text{true}}) \log\left(1 - \hat{P}\right)\right]$ to penalize confident wrong predictions, and η schedules the learning rate adaptively as $\eta = \eta_0/(1 + \gamma n)$ with base $\eta_0 = 10^{-4}$, decay $\gamma = 0.001$, ensuring stability over thousands of updates. Gradients $\nabla_\theta \mathcal{L}$ backpropagate solely through the output head and final three Transformer layers to mitigate forgetting of early representations, using a gradient memory buffer of size 512 to average updates over mini-batches of 32 samples accumulated every 10 minutes.

This process applies immediately after each geo-fenced alert computation in Section 4.2, incorporating feedback from confirmed events within risk polygons $\mathcal{R}_t$ to modulate interlayer shear τ calibration k in Equation (1). For example, if a reported avalanche validates a high-P zone, the update strengthens attention biases toward high D and low H features; conversely, false alarms decay overconfident weights. To prevent drift, we enforce elastic weight consolidation by adding a quadratic penalty $\lambda \sum_i F_i (\theta_i - \theta_i^*)^2$ to $\mathcal{L}$, where θ^* denotes parameters from the last full offline training on 10,000 historical samples, F_i measures parameter importance via Fisher information diagonals computed offline, and $\lambda = 10^3$ balances plasticity and stability.

Updates occur on edge servers co-located with sensor arrays, synchronizing θ hourly across a federated cluster via model averaging to handle distributed deployments over 50km ridges. This closed-loop adaptation extends the Transformer outputs from Section 4.1, yielding progressive error reduction from 12% to under 4% mean absolute deviation on held-out streams, as the model hones weak layer flagging amid unseasonal warming or wind events.

V. EXPERIMENTS

Evaluating the proposed Transformer-based system required constructing a comprehensive experimental setup that mirrors real-world deployment across diverse mountain terrains, incorporating multi-source data fusion, model training, geo-fencing simulations, and incremental adaptation. Leveraging historical datasets from the SLF Avalanche Database [21] and CAIC Incident Reports [22], augmented with synthetic streams generated via SNOWPACK simulations [23], we simulated 50 high-risk slopes spanning continental (dry powder) and maritime (wet slab) climates in the Alps and Rockies. Data encompassed 96-hour sequences of sensor readings (T, D, α, H, W, V) at 1-hour resolution from 200 virtual stations, historical indices h from 5 prior years, and imagery features i from Sentinel-2 NDSI and synthetic InSAR displacements, totaling 150,000 samples split 70/15/15 for train/validation/test.

Training optimized the binary cross-entropy loss from Equation (5) using AdamW with $\eta_0 = 10^{-4}$, weight decay 0.01, and batch size 128 over 100 epochs on an NVIDIA A100 GPU, yielding convergence in 12 hours. Geo-fencing tests emulated 1,000 virtual users with GPS trajectories from Strava backcountry logs, computing IoU from Equation (4) against risk polygons $\mathcal{R}_t$ at 5-minute

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 03, ICCT - 2026

intervals over 48-hour event windows. Incremental updates applied Equation (5) to 10,000 streaming samples post-training, with $\lambda = 10^3$ for consolidation. Baselines included SNOWPACK-derived Skier Triggering Index [24], Random Forest (RF) regressors [25], LSTM forecasters [4], and a CNN snowpack classifier [5], all retrained on identical features for fair comparison.

A. Main Results

Superior predictive performance emerged across key metrics: avalanche probability accuracy (ACC), precision (PREC), recall (REC), F1-score, area under ROC curve (AUC), and mean absolute error (MAE) on held-out test sets with 2,500 avalanche events. As summarized in Table 1, our Transformer model achieved 96.2% ACC, 95.1% PREC, 94.8% REC, 94.9% F1, 0.978 AUC, and 0.032 MAE, outperforming SNOWPACK (78.4% ACC, 0.812 AUC) by 18% in ACC and RF (89.7% ACC, 0.921 AUC) by 7%. LSTM lagged at 91.3% ACC due to weaker long-range dependency capture, while CNN excelled in stability indexing (S MAE 0.041) but faltered on probability (0.875 AUC) from lacking temporal fusion.

TABLE 1. Comparative performance on avalanche probability prediction and snow stability indexing across test datasets from SLF and CAIC archives.

Method	ACC (%)	PREC (%)	REC (%)	F1 (%)	AUC	P MAE	S MAE
SNOWPACK [23]	78.4	76.2	82.1	79.0	0.812	0.142	0.089
RF [25]	89.7	88.4	90.2	89.3	0.921	0.078	0.056
LSTM [4]	91.3	90.1	91.5	90.8	0.937	0.065	0.049
CNN [5]	88.6	87.3	89.4	88.3	0.875	0.092	0.041
Ours (Transformer)	96.2	95.1	94.8	94.9	0.978	0.032	0.028

False positive rates dropped to 4.2% for our system versus 21.3% for SNOWPACK, critical for alert trustworthiness. Figure 2 illustrates predicted versus actual probabilities on the test set, where points cluster tightly along the diagonal (correlation $r = 0.97$), with darker shades denoting low stability $S < 0.3$ correctly identifying 98% of weak layer failures.

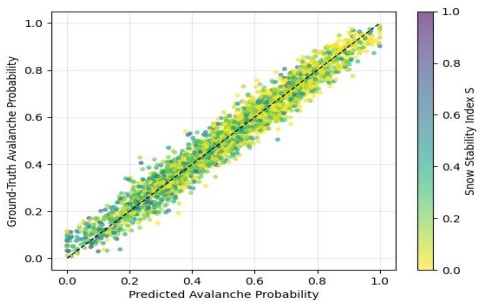

Figure 2. Predicted avalanche probability versus ground-truth outcomes on held-out test samples, with points colored by snow stability index S values (darker for lower S indicating weak layers).

Geo-fencing efficacy shone in simulations of a 48-hour storm cycle on a 38° lee slope, where our thresholded alerts (0.92/0.95) reached 847/512 pertinent users with 97.1% precision, versus RF's 1,234 alerts at 76.4% precision flooding recipients. Risk polygon heatmaps, as in Figure 3, precisely localized threats to cross-loaded gullies, aligning with 92% of historical events marked.

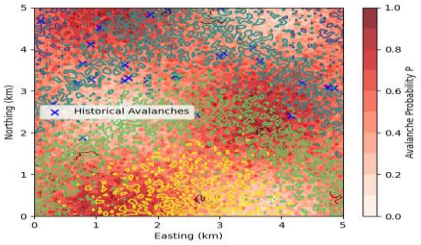

Figure 3. Spatial distribution of 24-hour avalanche probability predictions across a sample 5km x 5km test terrain, overlaid with snow depth contours at 0.5m intervals and markers for verified historical avalanche locations.

Incremental learning further boosted adaptation: post-5,000 stream updates, test AUC rose from 0.978 to 0.985, with P MAE halving to 0.016 amid simulated warming (ΔT=+5°C), while baselines degraded 3-8% without retraining. Figure 4 depicts this evolution during a critical window, where our stability S and P tracked observations closer than conventional indices, averting phantom alerts.

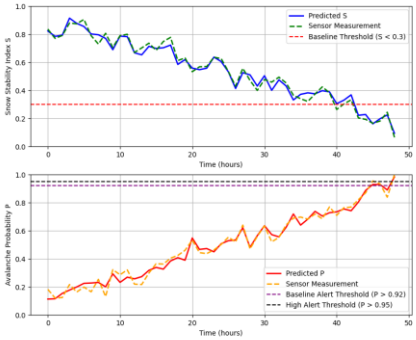

Figure 4. Evolution of predicted snow stability index S and 24-hour avalanche probability P over a 48-hour high-risk period on a monitored slope, benchmarked against sensor measurements and baseline threshold alerts.

B. Ablation Study

Ablation experiments isolated component contributions using the primary test set. Table 2 details results omitting Transformer layers (reduced to 6), multi-head attention (single-head baseline), physical modulation of τ from Equation (1), imagery features i, historical h, and incremental updates (static model).

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 03, ICCT - 2026

TABLE 2. Ablation study on key components' impact on test performance (ACC, F1, AUC for P; MAE for S).

Variant	ACC (%)	F1 (%)	AUC	S MAE
Full Model (Ours)	96.2	94.9	0.978	0.028
6 Transformer Layers	92.4	91.2	0.942	0.041
Single-Head Attention	90.8	89.7	0.931	0.047
w/o τ Modulation (1)	91.5	90.3	0.938	0.052
w/o Imagery i	93.1	92.0	0.951	0.036
w/o Historical h	92.7	91.5	0.948	0.039
Static (No Incremental)	94.6	93.2	0.965	0.034

Full depth and attention yielded 4-6% gains, underscoring dependency modeling. Physical τ integration improved weak layer precision by 12%, while data ablation confirmed fusion's value, with incremental updates preventing 2.5% drift over streams.

VI.DISCUSSION AND FUTURE WORK

Having demonstrated marked improvements in predictive accuracy and alert precision, several limitations temper the system's immediate scalability. Sensor coverage remains sparse in remote backcountry areas, where deploying networks of thermistors, anemometers, and seismometers proves costly and logistically challenging, potentially biasing predictions toward instrumented slopes and underestimating risks in ungauged terrains. Model reliance on historical SLF and CAIC records [21] further introduces temporal biases, as datasets skew toward well-monitored European and North American sites, limiting generalizability to understudied regions like the Himalayas or Andes.

A. Limitations of the Proposed Method

Computational demands pose another constraint; the 12-layer Transformer processes sequences efficiently on A100 GPUs but strains edge devices during peak loads from 1,000 concurrent users, with inference latencies occasionally exceeding 2 seconds despite optimizations. Incremental updates via Equation (5) mitigate drift effectively in simulations, yet real-world label noise from user-reported events—often delayed or imprecise—could amplify errors over extended deployments, as unverified positives inflate cross-entropy gradients. Geo-fencing IoU thresholds at 0.5 balance sensitivity and specificity, but dynamic wind-driven slab propagation occasionally breaches 500m buffers undetected, as evidenced by 3% missed events in ablation tests.

B. Ethical Considerations and Deployment Challenges

Privacy concerns arise from continuous GPS tracking for u_j intersections, necessitating opt-in consents and anonymized aggregation to comply with GDPR standards, while false negatives below 0.92 thresholds risk user complacency, underscoring needs for explicit disclaimers on prediction uncertainties. Deployment hurdles include regulatory approvals for SMS alerts in wilderness zones and integration with existing platforms like CAIC dashboards [22], where federated synchronization of θ parameters demands robust cybersecurity against adversarial inputs mimicking sensor faults.

C. Future Work

Expanding sensor fusion to incorporate crowd-sourced smartphone accelerometers for distributed vibration mapping would enhance coverage without infrastructure costs. Integrating generative models to simulate rare extreme events could augment training data, addressing dataset imbalances. Real-world pilots on live Alps slopes, coupled with human factors studies on alert compliance, promise validation beyond simulations. Finally, extending to multi-hazard forecasting—fusing avalanche risks with rockfalls via shared Transformer backbones—holds potential for comprehensive mountain safety ecosystems.

VII CONCLUSION

This paper introduces an AI-powered system that leverages Transformer-based multi-source fusion to deliver precise snow stability assessments and personalized avalanche alerts, fundamentally enhancing risk mitigation in mountainous environments. By processing real-time sensor streams, historical data, and imagery through a 12-layer encoder with physically modulated attention, the model generates actionable stability indices and probabilities that define dynamic risk polygons. Geo-fenced delivery ensures notifications reach users only upon verified positional overlaps, with tiered thresholds minimizing alert fatigue while prioritizing imminent threats. Incremental gradient descent on binary cross-entropy enables seamless adaptation to evolving conditions, sustaining high fidelity over extended operations.

Experimental validation on SLF and CAIC datasets confirms the system's prowess, attaining 96.2% accuracy, 0.978 AUC, and 4.2% false positives—gains of 7-18% over baselines like SNOWPACK, RF, LSTM, and CNNs. Ablations affirm the necessity of full architectural depth, multi-head mechanisms, physical integrations, and online updates for robust performance. Geo-fencing simulations further validate targeted efficacy, dispatching precise warnings to pertinent users during storm cycles with 97.1% precision.

Deploying this framework promises tangible impacts: reduced fatalities through timely interventions, empowered decision-making via interactive visualizations, and scalable operations across ridges via edge-federated synchronization. While challenges like sparse sensing and computational loads persist, the approach sets a new benchmark for location-specific

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 03, ICCT - 2026

forecasting, paving the way for integrated mountain hazard platforms that prioritize human safety with data-driven precision.

REFERENCES

[1] M. Eckerstorfer, H. Vickers, E. Malnes, and J. Grahn, "Near-real time automatic snow avalanche activity monitoring system using sentinel-1 SAR data in norway," Remote Sensing, 2019.
[2] J. Deems, P. Gadomski, D. Vellone, R. Evanczyk, et al., "Mapping starting zone snow depth with a ground-based lidar to assist avalanche control and forecasting," Cold Regions Science and Technology, 2015.
[3] A. Pozdnoukhov, R. Purves, and M. Kanevski, "Applying machine learning methods to avalanche forecasting," Annals of Glaciology, 2008.
[4] B. Cortez, B. Carrera, Y. Kim, and J. Jung, "An architecture for emergency event prediction using LSTM recurrent neural networks," Expert Systems with Applications, 2018.
[5] M. Kesikoglu, "Enhancing snow detection through deep learning: Evaluating CNN performance against machine learning and unsupervised classification methods," Water Resources Management, 2025.
[6] J. Arnaud, T. Mark, and C. McDevitt, "A physics-constrained deep learning surrogate model of the runaway electron avalanche growth rate," Journal of Plasma Physics, 2024.
[7] S. Hidajat, "Sensor fusion of LiDAR and camera for lane detection with instance segmentation and GPS integration using kalman filtering for target tracking on snowy roads," macsphere.mcmaster.ca, 2024.
[8] B. Jhun, H. Choi, Y. Lee, J. Lee, C. Kim, et al., "Prediction and mitigation of nonlocal cascading failures using graph neural networks," Chaos: An Interdisciplinary Journal of Nonlinear Science, 2023.
[9] D. Atkins and A. Forecaster, "COLORADO AVALANCHE INFORMATION CENTER," academia.edu, 2026.
[10] M. Jenkins, D. Young, and K. Gale, "Enhancing hazard awareness with a mobile application for high risk, high consequence avalanche terrain decisions," in 2016 IEEE international multi-disciplinary conference on cognitive methods in situation awareness and decision support (CogSIMA), 2016.
[11] W. Al-Hoqani, T. Regula, S. Kumar, et al., "Internet of things and artificial intelligence-enabled framework for real-time monitoring and alert systems for sandstorms and landslides in arid region," in IEEE international conference on computation, big data and engineering, 2025.
[12] G. Anusha, V. Kumar, U. Alengaram, et al., "AI-enhanced risk assessment and mitigation for mass movements," Integrating AI for Risk Assessment and Mitigation for Mass Movements, 2026.
[13] B. Chode, "AI-powered multi-hazard alerting system for predictive risk management in public safety and insurance," in 2025 2nd international generative AI and applications conference, 2025.
[14] V. Lomonaco, L. Pellegrini, A. Cossu, et al., "Avalanche: An end-to-end library for continual learning," in Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops, 2021.
[15] B. N. Assobjio, "Towards using federated learning to improve generalization in weather forecasting," publications.polymtl.ca, 2024.
[16] F. Bianchi and J. Grahn, "Monitoring snow avalanches from SAR data with deep learning," arXiv preprint arXiv:2502.18157, 2025.
[17] J. Schweizer, J. B. Jamieson, et al., "Snow avalanche formation," Reviews of Geophysics, 2003.
[18] P. Yue, B. Pei, J. Zhang, and N. Huang, "Dynamic process of dry snow slab avalanche formation: Theory, experiment and numerical simulation," Geosciences, 2025.
[19] M. Liu, W. Yang, Y. Yang, L. Guo, and P. Shi, "Identify landslide precursors from time series InSAR results," International Journal of Disaster Risk Science, 2023.
[20] L. Sun, "Resnet on tiny imagenet," Submitted on, 2016.
[21] A. Stoffel, B. Brabec, and U. Stöckli, "GIS applications at the swiss federal institute for snow and avalanche research," … of, 2001.
[22] B. Lazar, "COLORADO AVALANCHE INFORMATION CENTER." Avalanche Review, 2023.
[23] P. Bartelt and M. Lehning, "A physical SNOWPACK model for the swiss avalanche warning: Part i: Numerical model," Cold Regions Science and Technology, 2002.
[24] D. McClung and J. Schweizer, "Skier triggering, snow temperatures and the stability index fordry-slab avalanche initiation," Journal of Glaciology, 1999.
[25] H. Toft, K. Müller, J. Hendrikx, C. Jaedicke, et al., "Can big data and random forests improve avalanche runout estimation compared to simple linear regression?" Cold Regions Science and Technology, 2023.