Privacy-Preserving Multi-Modal Threat Detection Framework for Educational Facilities: A Conceptual Design

Abinaya Mettuapatti Sivagnanam
Independent Researcher
Texas, USA
https://orcid.org/0009-0004-8032-065X

Gayathri Surianarayanan
Independent Researcher
Virginia, USA
https://orcid.org/0009-0001-7987-5090

Srikanth Singireddy
Independent Researcher
North Carolina, USA
https://orcid.org/0009-0007-0528-0186

Vinoth Asoor Palanivel
Independent Researcher
Texas, USA
vinoth.itgeek@gmail.com

***Abstract*—School safety systems must detect threats rapidly while protecting student privacy and minimizing false alarms. This paper presents a conceptual framework combining edge-based open-vocabulary object detection with millimeter-wave radar sensing. Our three-tier confidence system matches response intensity to detection certainty, incorporating teacher verification at intermediate levels. The framework processes all data locally and includes automatic face anonymization, minimal retention policies, and no identity tracking. We provide practical camera placement methodology using pixels-per-meter calculations and detailed implementation guidelines. While this work presents a design proposal rather than experimental validation, it offers a comprehensive blueprint balancing security effectiveness with ethical requirements. The framework is designed to be reproducible and adaptable to diverse educational environments.**

***Index Terms*—Conceptual framework, edge computing, open-vocabulary detection, mmWave radar, privacy-preserving systems, human-in-the-loop, school security**

I. Introduction

Educational institutions need safety systems that detect potential threats early while respecting student privacy and maintaining normal learning environments. Current approaches have significant limitations: metal detectors create bottlenecks and miss non-metallic threats, traditional AI systems require expensive retraining when threat definitions change, and cloud-based solutions raise privacy concerns while introducing network latency.

Recent advances in open-vocabulary object detection [1] and millimeter-wave radar [2] create new possibilities. Open-vocabulary detectors recognize objects from text descriptions without prior training on specific classes. mmWave radar has demonstrated potential for detecting concealed rigid objects through fabric and certain non-metallic materials in controlled laboratory conditions [2], though real-world performance varies significantly with material composition, object size, range, and environmental factors. Modern edge processors can run these algorithms locally with sub-second latency. However, no comprehensive framework exists combining these technologies while addressing both technical and ethical requirements.

A. Contributions

This conceptual framework paper provides:

1) System architecture combining open-vocabulary vision with mmWave radar in an edge-computing framework (Section III).
2) Three-tier confidence system matching response intensity to detection certainty with selective human verification (Section IV).
3) Privacy-by-design principles including edge processing, automatic anonymization, and minimal data retention (Section V).
4) Deployment methodology with camera placement algorithms based on pixels-per-meter calculations (Section VI).
5) Implementation roadmap with technology recommendations and validation strategies (Section VII).
6) Design targets and a simple validation plan that others can use to run pilots later.

This is a design proposal rather than experimental validation. We provide a blueprint for future research and deployment that others can reproduce and test.

II. Background

A. Open-Vocabulary Detection

Traditional security systems use fixed-category detectors trained on specific threats [3]. These require retraining for new categories. Open-vocabulary models like Grounding DINO 1.5 [1] and OWL-ViT [4] learn to ground natural language in visual features, enabling zero-shot detection. Administrators can describe threats with text prompts like ”handgun” or ”long rigid object” without labeled training examples.

B. Millimeter-Wave Sensing

mmWave radar operates at 60-81 GHz and has demonstrated the ability to penetrate fabric and some plastics while reflecting from metal and rigid objects in controlled experimental settings, though detection performance varies substantially with material composition, object geometry, standoff distance, and clutter conditions [2]. FMCW radar provides range and velocity information through Doppler analysis. Recent work demonstrates potential effectiveness for concealed object detection under laboratory conditions [2]. Compact commercial modules make this technology practical for integration into security systems, though operational performance in diverse real-world environments requires systematic validation. Our framework positions radar as a complementary cue to vision rather than a standalone detection modality.

C. Multi-Modal Fusion and Human-in-the-Loop

Vision and radar provide complementary information. Vision excels at shape recognition; radar penetrates barriers and works in poor lighting. Sensor fusion improves robustness [5]. Human-in-the-loop designs balance automation with judgment [6], with AI highlighting cases for expert review [7].

D. Privacy in Educational Settings

Privacy concerns limit surveillance adoption in schools [8]. Traditional anonymization can be reversed [9]. Recent generative techniques provide stronger guarantees [10]. Edge computing keeps sensitive data local [11]. Our framework combines edge processing, anonymization, and strict retention policies.

III. Proposed System Architecture

Figure 1 shows the system architecture. All processing occurs on local edge devices. Video and radar data are captured simultaneously and time-synchronized. A processing pipeline performs detection, segmentation, tracking, and fusion. A confidence classifier routes events to three response tiers.

A. Edge Computing Foundation

The system runs on local edge hardware for privacy (raw video never leaves premises), latency (no network round-trip), and reliability (operates during internet outages). Edge devices include AI accelerators (GPU/TPU) for real-time processing.

B. Open-Vocabulary Detection

We propose Grounding DINO 1.5 [1] as the primary detector. Administrators configure threat categories using text:

Positive: ”handgun”, ”knife”, ”rifle”, ”long rigid metallic object”

Negative: ”umbrella”, ”tripod”, ”musical instrument case”

The model computes similarity between visual features and text embeddings, generating bounding boxes and confidence scores. This enables instant policy updates without retraining.

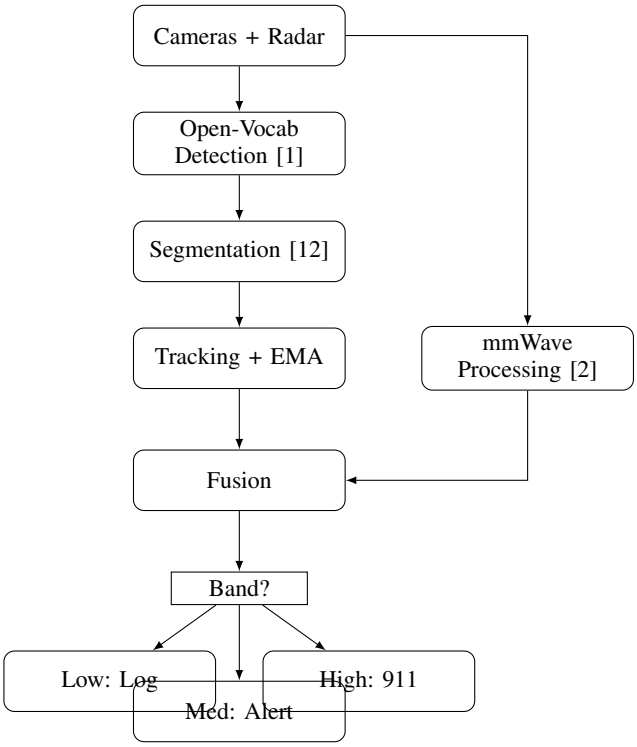

Fig. 1. System architecture: sensors through confidence-based response routing.

C. Segmentation and Tracking

Bounding boxes often include background clutter. We propose SAM2 [12] for precise object masks. Compute IoU between mask and box:

$$\text{IoU} = \frac{\text{Area}(\text{mask} \cap \text{box})}{\text{Area}(\text{mask} \cup \text{box})} \tag{1}$$

Low IoU (< 0.6) indicates loose boxes likely capturing clutter. Short-term tracking with exponential moving average (EMA) smooths confidence scores:

$$s_t = \alpha \cdot s_{\text{current}} + (1 - \alpha) \cdot s_{t-1} \tag{2}$$

where $\alpha \in [0.3, 0.4]$. This suppresses single-frame spikes.

D. mmWave Radar and Fusion

A compact FMCW radar [2] extracts range, velocity, and reflectivity. For each tracked object, we define a spatial window in radar coordinates and extract returns. Strong, stable returns from dense materials support threat hypothesis.

Late fusion combines scores:

$$s_{\text{fused}} = w_v \cdot s_{\text{vision}} + w_r \cdot s_{\text{radar}} \tag{3}$$

Typical weights: $w_v = 0.7$, $w_r = 0.3$, tuned during calibration.

E. Confidence Classification

Fused scores map to three bands using thresholds $(t_{\text{low}}, t_{\text{high}})$:

- **Low** ($s < t_{\text{low}}$): Log only
- **Medium** ($t_{\text{low}} \leq s < t_{\text{high}}$): Teacher verification
- **High** ($s \geq t_{\text{high}}$): Emergency response

IV. Response Protocol

A. Low Confidence: Passive Logging

Events below t_{low} are logged with timestamp, location, and anonymized thumbnail but trigger no alerts. Logs support pattern analysis and system tuning.

B. Medium Confidence: Silent Teacher Alert

Medium-band events initiate human verification before escalation, preventing false alarms while ensuring genuine threats receive response.

The system identifies teachers in the affected zone or adjacent zones. Alerts deliver via local network with:

- Zone name, timestamp
- Short anonymized video clip (1-3 seconds)
- Action buttons: "Confirm Threat" / "False Alarm"
- Optional notes field

All faces except the subject are automatically blurred [9], [10]. Proposed timeout: 20 seconds. Early termination if teacher responds or confidence rises to High. Teacher confirmations escalate to emergency response; dismissals drop to Low.

C. High Confidence: Emergency Response

High-band events trigger:

- **Zoned alarms**: Audio alerts in affected zones only
- **Door sequencing**: Controlled access maintaining free egress per NFPA 101 [13] (panic bars, fire/power auto-release)
- **NG911 notification**: Structured data to PSAP [14] with location, timestamp, threat description, anonymized images

A central console operator can view events, pause responses, override classifications, and manually initiate protocols.

V. Privacy and Ethics

A. Privacy-by-Design

Edge processing: Raw video never leaves premises.

Automatic anonymization: All shared media undergoes face anonymization using privacy-preserving techniques [10] that resist reconstruction [9].

Minimal retention: Alert clips 24-72 hours, anonymized logs 30 days, original video 7-14 days encrypted locally.

No identification: No facial recognition, watchlists, or identity linking.

B. Bias and Fairness

Test performance across diverse contexts (lighting, crowding, angles, distances) rather than demographic labels. Anonymization applies uniformly. All thresholds are documented and adjustable with logged justification.

C. Safety Compliance

Door automation follows NFPA 101 [13]: free exit always, automatic release on fire/power failure, manual override, audit logging. NG911 follows established protocols [14].

D. Governance

Roles: Console operator (monitoring), principal (policy approval), safety officer (threshold management), teachers (verification), IT (maintenance).

Change control: Modifications to prompts or thresholds require proposal, testing, approval, version control, and documentation.

Regular review: Quarterly examination of alert frequency, response times, system performance, and privacy compliance.

E. Threats and Failure Modes

Table I lists common risks and mitigation strategies.

F. Risk Matrix

Table II rates risks by likelihood and impact.

VI. Conceptual Design: Targets & Future Validation Plan

This section defines what "good" looks like for this system. These are targets and design-time estimates, not measured results.

A. Design Targets

Table III presents conceptual design targets.

B. Design-Time Estimates

Latency budget. We sum stage estimates from device and software specifications:

$$L^{95}_{\text{total}} = L^{95}_{\text{cap}} + L^{95}_{\text{vision}} + L^{95}_{\text{radar}} + L^{95}_{\text{fuse}} + L^{95}_{\text{net}} + L^{95}_{\text{ui}} \tag{4}$$

These values come from vendor documentation and will be replaced by measured traces during pilots.

False positives per hour. For frame rate F frames/hour/camera, per-frame vision false positive p_v, fraction that pass fusion q, and fraction escalated after teacher review r:

$$\text{FP/h/cam} \approx F \cdot p_v \cdot q \cdot r \tag{5}$$

Coverage vs. PPM. We report the percent of floor area that meets a chosen PPM target from camera geometry.

C. Pre-Registered Scenarios

Table IV lists validation scenarios with acceptance criteria.

TABLE I
THREATS AND FAILURE MODES — CONCEPTUAL

Mode	Detection Risk	Mitigation	Fail-Safe
Toy gun	May look like real weapon in video. Radar return is small or missing.	Use fusion to require both vision confidence and expected radar strength. Add teacher quick check before high tier.	Do not lock exits. Keep people safe. Notify staff quietly and watch the clip.
3D printed replica	Visual looks real. Plastic has weak radar return.	Lower confidence if texture and material cues disagree with radar. Escalate to medium tier for human review.	Keep doors in safe egress state. Alert resource officer and log for follow up.
Umbrella or tripod	Shape can trigger false alarms on vision.	SAM style mask and box overlap check. Context rules for posture and motion. Fusion lowers score if radar signature does not match metal.	Rate limit repeat alerts in the same scene. One tap dismiss keeps it quiet for a short window.
Camera outage	Blind spot and missed events.	Overlap coverage in layouts. Device heartbeat and health checks. Instant ticket to ops on loss.	Default to safety. Keep egress clear. Send hall monitor or security to the zone.
Radar ghost	Multipath or clutter creates phantom blobs.	Static clutter map, Doppler and track persistence filter. Require short visual trace before escalate.	Treat as low risk. Do not escalate to high tier without visual support.
Network jitter	Delayed or bursty alerts.	Local buffer and retry with backoff. QoS marking on alert channel. Show local on-device banner if delay grows.	If delay crosses a limit, notify staff locally and fall back to offline logging until link is stable.
Adversarial spoof (stickers/printed decoys)	Visual tricking of the model or prompts designed to confuse it.	Radar/material cues may disagree. Sanity checks for texture/material, mask–box overlap tests, and simple "commonsense" rules. Always require a quick teacher check before high tier.	Treat as medium tier. Never lock egress. Keep people safe, log the case, and improve filters in the next update.

TABLE II
LIKELIHOOD × IMPACT RISK MATRIX

Impact \ Likelihood	Low	Medium	High
High	Toy gun	3D-print replica; adv. spoof	
Medium	Camera outage	Network jitter	
Low	Umbrella/tripod	Radar ghost	

TABLE III
DESIGN TARGETS (CONCEPTUAL, NOT MEASURED)

Metric	Target
Detection recall (threat objects)	≥ 0.80
False positives (per camera per hour)	≤ 0.05
End-to-end alert latency (P95)	≤ 30 s
Teacher verification time (P95)	≤ 20 s
Coverage at PPM threshold	$\geq 95\%$
Availability during school hours	$\geq 99.5\%$

TABLE IV
SCRIPTED SCENARIOS AND ACCEPTANCE CRITERIA (FOR FUTURE PILOTS)

ID	Description	Success Metric	Target
S1	Umbrella or tripod (non-threat)	Alerts/hour/camera	≤ 0.05
S2	Concealed toy handgun	Detection recall	≥ 0.80
S3	Camera occlusion (30 s)	Availability	$\geq 99\%$
S4	Network jitter (200 ms)	Alert P95 latency	≤ 30 s
S5	Busy hallway clip review	Teacher verify P95	≤ 20 s
S6	Privacy audit	Stored PII	0 items

D. Calibration and Ablation

We will choose thresholds $(t_{\text{low}}, t_{\text{high}})$ to meet the false-positive target while keeping recall high; fuse vision/radar with (w_v, w_r); compare (vision-only), (radar-only), and (fused) pipelines; and measure human-in-the-loop timing on anonymized clips. This manuscript defines evaluation scenarios, metrics, and acceptance criteria; empirical results are intentionally out of scope and are reserved for a follow-on study after a controlled pilot.

VII. DEPLOYMENT METHODOLOGY

A. Camera Placement

Effective detection requires proper placement. We provide systematic guidance based on geometric analysis.

1) Coverage by Zone: **Entries**: Cover approach and doorway; co-locate radar; eliminate window glare.

Corridors: Maintain 20-25% overlap between adjacent cameras.

Common areas: Multiple angles to reduce occlusion; overhead camera near busy doors.

2) Pixels-Per-Meter Calculation: For camera with horizontal resolution N_x and field of view θ, scene width at distance d is $W(d) = 2d\tan(\theta/2)$. Pixels per meter:

$$\text{PPM}(d) = \frac{N_x}{2d\tan(\theta/2)} \tag{6}$$

To ensure minimum detection quality (e.g., $\text{PPM}_{\text{target}} = 25$ for handgun-sized objects), maximum effective distance:

$$d_{\max} = \frac{N_x}{2 \cdot \text{PPM}_{\text{target}} \cdot \tan(\theta/2)} \tag{7}$$

Camera spacing: $\leq 0.75 \cdot d_{\text{max}}$ to maintain overlap.

Example: For 1080p ($N_x = 1920$), 90° FOV, $\text{PPM}_{\text{target}} = 25$:

$$d_{\text{max}} = \frac{1920}{2 \times 25 \times \tan(45)} = 38.4 \text{ m} \quad (8)$$

Spacing: ≤ 28.8 m. For 4K ($N_x = 3840$): $d_{\text{max}} = 76.8$ m, spacing ≤ 57.6 m.

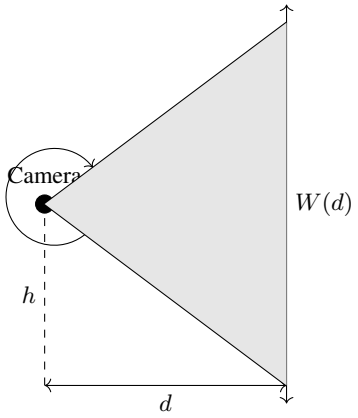

$$W(d) = 2d\tan(\theta/2)$$
$$\text{PPM}(d) = N_x/W(d)$$

Fig. 2. Camera coverage geometry: field of view θ determines scene width $W(d)$ at distance d, which determines pixels-per-meter for detection quality.

3) Installation Guidelines: Height 3-4 m, tilt 15-25°. Use corridor lenses (narrow FOV) for halls, wide lenses for open areas. Avoid backlight and obstructions.

B. System Configuration

Table V summarizes key parameters. Actual values determined during implementation.

TABLE V
PROPOSED CONFIGURATION PARAMETERS

Parameter	Proposed Value
Detection model	Grounding DINO 1.5 [1]
Segmentation model	SAM2 [12]
Mask-box IoU threshold	0.60
Track duration	0.5-1.0 seconds
EMA factor	0.3-0.4
Radar frequency	60-81 GHz FMCW [2]
Vision weight (w_v)	0.6-0.8 (tune)
Radar weight (w_r)	0.2-0.4 (tune)
Low threshold	0.4-0.5 (tune)
High threshold	0.7-0.8 (tune)
Teacher timeout	15-30 seconds

VIII. IMPLEMENTATION ROADMAP

A. Technology Selection

Vision: Grounding DINO 1.5 or OWL-ViT; criteria: zero-shot performance, prompt flexibility, inference speed.

Segmentation: SAM2; criteria: prompt-based operation, temporal consistency.

Edge platform: NVIDIA Jetson (strong AI), Intel GPU servers (flexible), or specialized accelerators (efficient). Selection based on camera count, resolution, latency requirements.

Radar: Texas Instruments AWR series, Infineon 60 GHz, or NXP automotive. Selection based on range/angular resolution, SDK availability.

B. Implementation Phases

Phase 1 (2-3 months): Site survey, camera placement design, hardware selection, privacy policy development, community engagement.

Phase 2 (2-3 months): Pilot installation in limited area, basic detection pipeline, initial data collection.

Phase 3 (3-4 months): Integrate detection models, segmentation, tracking, radar processing, fusion logic, anonymization.

Phase 4 (2-3 months): Collect validation scenarios, tune thresholds, conduct drills, measure latency, refine interface.

Phase 5 (3-6 months): Full deployment, NG911 integration, door control integration, staff training, full drills.

Phase 6 (ongoing): Performance monitoring, quarterly review, annual audit, continuous training.

C. Validation Strategy

Since this is a conceptual framework, we outline how implementers should validate:

Synthetic testing: Use public datasets to test open-vocabulary detection on threat-like objects. Establish baseline performance.

Staged scenarios: Controlled tests with known objects. Vary lighting, distance, crowding, angles. Test false positive sources (umbrellas, tripods). Measure accuracy, latency.

Drill exercises: Full-system tests with law enforcement. Gather timing data. Collect teacher feedback. Identify procedure gaps.

Key metrics: Precision, recall, F1, false positive rate (per hour), false negative rate, frame-to-detection latency (median/P95), detection-to-alert latency, teacher response time, system availability.

Design goals (hypotheses to be validated in pilots): detection precision ≥ 0.80, recall ≥ 0.75, false positives ≤ 0.1/hour/camera, median frame-to-detection ≤ 1 s, and teacher response P95 ≤ 30 s. These targets are design goals derived from component literature [1], [2], [12] and typical edge pipeline budgets; they require empirical validation in school-like conditions including variable lighting, occlusion, and motion blur.

D. Cost Considerations

Approximate ranges:

- IP camera $200-800
- Edge compute node $1,000-5,000
- Radar module $500-2,000
- Small facility (10 cameras): $20,000-40,000
- Medium (30 cameras): $50,000-100,000

Costs reducible through existing infrastructure, open-source software, phased deployment.

IX. Discussion

A. Advantages

Technical: Adaptability via text prompts, complementary vision-radar sensing, edge processing eliminates cloud latency, graduated response balances speed and caution.

Operational: Human oversight filters false positives, transparent prompts enable community trust, modular architecture supports customization, code-compliant door logic.

Ethical: Multiple privacy layers, no identification, minimal retention, community transparency.

B. Limitations

Technical: Performance depends on lighting and viewing angles. Distance limits based on resolution. Novel threat presentations may not match prompts. Significant computational requirements.

Operational: Threshold calibration requires extensive testing. Frequent alerts could cause teacher fatigue. Regular maintenance overhead. Integration complexity with existing systems.

Ethical: Privacy perceptions vary. Cannot prevent all incidents. Cost disparities between schools. Potential psychological effects on students.

C. Suitable Contexts

Most appropriate for schools with technical resources, communities prioritizing security and privacy, facilities with existing IP infrastructure, and institutions committed to transparent governance. Less suitable for very small schools, resource-constrained districts, or communities opposing any surveillance.

X. Future Research Directions

Technical improvements: Learned attention fusion, behavioral analysis (with ethical care), active learning from teacher feedback, explainability enhancements.

Evaluation: Large-scale validation across diverse schools, long-term performance studies, human factors research on teacher workload and student perceptions.

Policy: Model privacy frameworks, legal compliance guidelines, ethical standards, comparative effectiveness studies.

XI. Conclusion

We have presented a conceptual framework for multi-modal threat detection in schools that balances security, privacy, and practicality. The system combines open-vocabulary detection with mmWave radar, processes data on edge devices, and implements graduated response through three-tier confidence classification. Key principles include privacy-by-design, adaptive text-based configuration, selective human verification, safety compliance, and practical deployment methodology.

No technology can eliminate all threats. Deployment must be part of comprehensive safety strategies including prevention, training, emergency preparedness, and mental health support. Success requires technical excellence, institutional commitment, community trust, and ongoing attention to ethical implications.

We hope this framework stimulates productive discussion and provides practical guidance for researchers, practitioners, and policymakers working to improve school safety while respecting privacy, equity, and human dignity.

Ethics and Consent Statement

Deployment of surveillance systems in educational settings requires careful ethical consideration and community engagement. Any implementation of this framework must include: (1) transparent communication with students, parents, staff, and community stakeholders about system capabilities and limitations; (2) formal governance structures with clear policies for system operation, data access, and audit procedures; (3) documented consent processes respecting local regulations and institutional policies; (4) mechanisms for opt-out or accommodation where feasible; and (5) institutional review board (IRB) or ethics committee oversight for any pilot studies involving human subjects. The authors acknowledge that technical feasibility does not imply ethical or legal permissibility, and strongly recommend consultation with legal counsel, privacy experts, and community representatives before any deployment.

Acknowledgments

We used AI for grammar and language polishing. All ideas are our own.

References

[1] T. Ren et al., "Grounding DINO 1.5: Advance the 'edge' of open-set object detection," *arXiv:2405.10300*, May 2024. [Online]. Available: https://arxiv.org/abs/2405.10300. DOI: 10.48550/arXiv.2405.10300.

[2] C. Kaul et al., "AI-enabled sensor fusion of time-of-flight imaging and mmWave for concealed metal object detection," *Sensors*, vol. 24, no. 18, p. 5865, Sept. 2024. DOI: 10.3390/s24185865. [Online]. Available: https://www.mdpi.com/1424-8220/24/18/5865.

[3] G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," Software v8.0.0, Jan. 2023. [Online]. Available: https://github.com/ultralytics/ultralytics. [Accessed: Jan. 2026].

[4] M. Minderer et al., "Simple open-vocabulary object detection with vision transformers," in *Proc. ECCV*, 2023, pp. 728-755. DOI: 10.1007/978-3-031-20077-9_42.

[5] B. Khaleghi et al., "Multisensor data fusion: A review of the state-of-the-art," *Information Fusion*, vol. 14, no. 1, pp. 28-44, Jan. 2013. DOI: 10.1016/j.inffus.2011.08.001.

[6] R. Monarch, *Human-in-the-Loop Machine Learning*. Manning, 2023. ISBN: 978-1617296741.

[7] A. Chen, J. Liu, and M. Zhang, "Human-AI collaboration in medical image analysis: A survey," *Medical Image Analysis*, vol. 92, p. 102951, Feb. 2024. DOI: 10.1016/j.media.2023.102951.

[8] K. Greene and T. Johnson, "Privacy concerns in educational surveillance technologies," *Educational Technology & Society*, vol. 26, no. 2, pp. 145-158, Apr. 2023.

[9] J. Todt et al., "Fantômas: Understanding face anonymization reversibility," in *Proc. PETS*, vol. 2024, no. 4, pp. 654-671, 2024. DOI: 10.56553/popets-2024-0121.

[10] S. Park et al., "Privacy-driven faces: A survey on generative facial de-identification," *ACM Computing Surveys*, vol. 57, no. 8, pp. 1-38, 2025. DOI: 10.1145/3689028.

[11] L. Wang, S. Kumar, and R. Patel, "Privacy-preserving edge computing for smart surveillance," *IEEE Internet of Things Journal*, vol. 11, no. 8, pp. 13421-13435, Apr. 2024. DOI: 10.1109/JIOT.2023.3334241.

[12] N. Ravi et al., “SAM 2: Segment anything in images and videos,” *arXiv:2408.00714*, Aug. 2024. [Online]. Available: https://arxiv.org/abs/2408.00714. DOI: 10.48550/arXiv.2408.00714.

[13] National Fire Protection Association, “NFPA 101: Life Safety Code, 2024 Edition,” Quincy, MA, USA: NFPA, 2024. [Online]. Available: https://www.nfpa.org/codes-and-standards/1/0/1.

[14] National Emergency Number Association (NENA), “Security for Next Generation 9-1-1, NENA-STA-040.2-2024,” Arlington, VA, USA: NENA, Nov. 2024. [Online]. Available: https://www.nena.org/page/NG911SecurityTechStd.