

Seeing Danger Before Moving: Learning Environment-Centric Risk for Safe Robot Navigation

Palash Yuvraj Ingle* and Young-Gab Kim^{†‡}

*Department of Computer Information and Security
Sejong University, Seoul, South Korea
Email: palash@sejong.ac.kr

[†]Department of Computer and Information Security

[‡]Department of Convergence Engineering for Intelligent Drone
Sejong University, Seoul, South Korea
Email: alwaysgabi@sejong.ac.kr

[‡]Corresponding author: Young-Gab Kim

Abstract—Hazardous environments are characterized by uncertain, rapidly changing conditions, where autonomous robots are responsible for making safety-critical navigation decisions to ensure a safer path. When navigation relies solely on reactive perception, a robot may encounter dangerous situations before it has sufficient time to respond effectively. In real-world settings, however, hazards such as smoke, heat, flooding, or structural instability can be observed directly by the environment itself through cameras and sensors. This creates an opportunity to anticipate risk in advance, before a robot enters unsafe regions. In this work, we focus on enabling proactive and risk-aware robot navigation through environment-centric perception. We present a unified multimodal spatiotemporal model that estimates navigation risk by fusing environmental visual and sensory data through early integration, producing instantaneous probabilistic risk estimates. To avoid reacting to short-lived disturbances and sensor noise, these estimates are incorporated into a Bayesian filtering process that treats risk as an underlying latent state and maintains a stable belief over evolving environmental hazards. The resulting risk belief is integrated into the robot’s navigation stack as a cost term, guiding the planner toward safer routes and allowing the robot to avoid hazardous areas before reaching them. We evaluate the proposed approach on a physical multi-hazard testbed that includes smoke, heat, water presence, vibration, and structural obstacles. Experimental results show reduced hazard exposure and safer navigation behavior compared to risk-unaware baselines, highlighting the benefits of environment-centric risk inference for autonomous robot navigation.

I. INTRODUCTION

Robots are increasingly deployed in safety-critical settings such as disaster response, industrial incidents, and damaged infrastructure, where environmental conditions can change rapidly and hazards may be difficult for humans to assess directly. In these domains, navigation decisions are not merely about reaching a goal efficiently, but about minimizing exposure to dangerous operating conditions (e.g., smoke, heat, gas, flooding, or structural instability) that can endanger the robot, the mission, or nearby people [12].

A dominant paradigm in mobile robotics is *robot-centric*

navigation, where a robot perceives the world using onboard sensors, builds a local representation, and reacts as hazards become observable. Risk-aware navigation has advanced substantially in recent years, including pipelines that explicitly model risk during mapping and planning in challenging terrain [13] and classical formulations that enforce probabilistic safety constraints in motion planning [2]. More broadly, planning under partial observability is often posed in belief space, where decisions are based on a maintained belief over latent state [10]. However, in hazardous environments, purely reactive or robot-only perception can be fundamentally limited: a robot may encounter a dangerous region before it has sufficient time, visibility, or sensing reliability to respond conservatively.

In many real facilities, hazards are not only observable from the robot itself. The *environment* may already be instrumented with cameras and heterogeneous sensors, or such infrastructure can be rapidly deployed, providing measurements that extend beyond the robot’s perceptual horizon. Prior work on integrating robots with sensor networks highlights this complementarity: distributed sensing can provide early situational awareness, while robots execute actions and close the loop in the physical world [5]. Related perspectives in cloud and networked robotics further emphasize how external infrastructure can augment robot capabilities through shared sensing and computation [11]. These observations motivate an alternative viewpoint: *environment-centric risk perception*, where navigation risk is inferred primarily from infrastructure sensing and then injected into the robot’s navigation stack.

Despite this potential, existing risk-aware navigation systems largely treat environmental risk as an instantaneous signal or a static cost derived directly from raw observations. Such formulations are sensitive to sensor noise, partial coverage, and transient disturbances, often leading to unstable planning behavior or overly conservative avoidance. What is missing is a principled way to represent environmental risk as a

temporally persistent latent state that can be reliably used for planning under uncertainty.

In this paper, we address this gap by introducing an environment-centric risk-aware navigation framework that explicitly models navigation risk as a latent, time-varying belief. Environment-mounted cameras and sensors are fused through a unified multimodal encoder to produce instantaneous probabilistic risk observations. Rather than directly reacting to these noisy estimates, we treat them as evidence and maintain a *filtered latent risk belief* using Bayesian state estimation principles [14]. This belief captures hazard persistence and uncertainty, and is converted into a planning-compatible risk costmap that can be composed with standard navigation costs. A planner (e.g., A* [7]) then selects routes that minimize expected hazard exposure. Importantly, the robot itself performs no hazard perception and relies on onboard sensing only for localization and motion execution.

We make the following contributions: (i) a formulation of environment-centric navigation risk as a latent, temporally evolving belief suitable for planning under uncertainty; (ii) a multimodal perception pipeline that fuses environment-mounted visual and heterogeneous sensor data to produce probabilistic risk observations; (iii) a risk costmap construction that integrates the latent belief with standard robot navigation planners without modification; and (iv) real-world validation in a physical multi-hazard testbed demonstrating proactive hazard avoidance and reduced exposure compared to risk-unaware and reactive baselines.

An overview of the proposed framework (*EnvGuard*) is shown in Fig. 1. The project page and source code for this work are available at: <https://github.com/palashngl/EnvGuard>.git.

The remainder of this paper is organized as follows. Section II reviews related work in disaster robotics, risk-aware navigation, probabilistic planning, and infrastructure-assisted sensing. Section III formulates environment-centric navigation risk and defines the latent risk belief and planning objective. Section IV presents the proposed perception, Bayesian filtering, and risk-aware planning framework. Experimental validation in a physical multi-hazard testbed is provided in Section V, followed by conclusions in Section VI.

II. RELATED WORK

a) Disaster robotics and hazard-aware autonomy:

Robotic operation in disaster environments has long been recognized as safety-critical due to degraded visibility, unstructured terrain, and time-varying hazards such as smoke, fire, flooding, and structural instability. Murphy provides a comprehensive account of real-world disaster robotics deployments and highlights the persistent gap between laboratory autonomy and fielded systems, emphasizing that reliable sensing and actionable situational awareness remain central bottlenecks for safe response operations [12]. Motivated by these challenges, a large body of work has focused on improving robot-side perception, robustness, and autonomy under adverse conditions. However, in many real deployments, hazards can also be

observed *externally* through fixed cameras and environmental sensors deployed as part of the infrastructure. Despite this availability, such environment-centric sensing is rarely exploited to explicitly infer and anticipate navigation risk before a robot physically enters hazardous regions.

b) *Risk-aware planning under uncertainty*: Planning and control under uncertainty have been extensively studied in robotics. Chance-constrained motion planning enforces probabilistic safety margins to account for uncertainty in perception, actuation, and environment dynamics [2]. At the control level, safety mechanisms such as control barrier functions (CBFs) provide formal guarantees by maintaining forward invariance of safe sets [1]. In navigation, many risk-aware pipelines derive costs from robot-centric observations such as terrain geometry, traction estimates, or traversability uncertainty, and then optimize motion to trade efficiency against safety. For example, RAMP constructs risk-aware maps and planners for fast off-road navigation by explicitly reasoning about unknown space and risk constraints [13]. EVORA learns aleatoric and epistemic uncertainty in traversability estimation and integrates this uncertainty into risk-aware planning for off-road autonomy [3]. While these approaches demonstrate the importance of uncertainty-aware costs, they primarily estimate risk from *robot-side* sensing and interaction. As a result, hazards are often detected only after the robot approaches or enters unsafe regions. In contrast, our work focuses on environment-centric risk inference, where hazards are estimated externally and stabilized over time before being incorporated into the navigation stack, enabling proactive rather than reactive decision making.

c) *Belief-space estimation and probabilistic maps*: Our formulation is grounded in probabilistic state estimation and belief-space reasoning. Bayesian filtering provides the standard framework for recursively estimating latent states from noisy observations [14]. More generally, POMDP formulations formalize decision making under partial observability by planning over belief states rather than deterministic state estimates [10]. Classic probabilistic mapping approaches, such as occupancy grids, treat occupancy as a latent variable inferred from sensor evidence accumulated over time [6]. We adopt a similar philosophy, but estimate a latent *environmental risk* state rather than occupancy. Importantly, our belief does not represent robot state uncertainty; instead, it captures the evolving hazard level of the environment itself, which is subsequently projected into a risk costmap used by a standard navigation planner.

d) *Infrastructure-assisted and networked sensing for robotics*: Leveraging infrastructure to extend a robot's perceptual horizon has strong precedent in networked robotics and wireless sensor network (WSN) research. Environmental WSNs have been widely studied for long-term monitoring using heterogeneous sensing modalities over large areas [4]. Frameworks integrating robots with sensor networks demonstrate how external sensing can complement onboard perception and support cooperative tasks [5]. Surveys on mobile robots in WSN settings categorize common task structures




Fig. 1: Overview of the proposed environment-centric risk-aware navigation framework (*EnvGuard*). Hazard perception is performed exclusively by environment-mounted cameras and sensors, which estimate instantaneous risk observations that are fused and filtered into a latent risk belief representing the evolving hazard state of the environment. This belief is transformed into a risk costmap and injected into a standard planner to compute safe navigation paths. The robot does not perform hazard perception and relies solely on onboard perception for localization and motion execution.

such as monitoring, coverage, data collection, and servicing [8]. Cloud robotics further motivates off-board computation and shared perception resources when onboard sensing or compute is limited [11]. Despite this progress, relatively few works treat infrastructure-mounted sensing as a first-class input to a navigation-risk inference model that explicitly produces a temporally filtered hazard belief and couples it to a standard navigation stack. Our work addresses this gap by learning environment-centric multimodal risk, stabilizing it through Bayesian belief filtering, and validating its impact on navigation through real-world experiments in a multi-hazard physical testbed.

e) Summary and positioning: In summary, prior work on risk-aware navigation largely focuses on robot-centric perception, instantaneous uncertainty estimation, or reactive safety mechanisms. In contrast, our approach treats environmental hazard as a latent, temporally evolving state inferred from infrastructure-mounted sensing. By combining environment-centric multimodal perception with Bayesian belief filtering and direct integration into an existing navigation stack, our

method enables proactive risk-aware navigation that anticipates hazards before physical exposure. This distinction is validated through real-world experiments in a physical multi-hazard testbed, which remains underexplored in existing risk-aware navigation literature.

III. PROBLEM FORMULATION

This section formalizes environment-centric navigation risk and defines the latent risk belief, observation model, and planning objective that underpin the proposed framework.

a) Environment-centric sensing: We consider a planar workspace $\Omega \subset \mathbb{R}^2$ instrumented with *environment-mounted* cameras and heterogeneous sensors. At time t , the environment provides (i) visual observations from M fixed cameras, $\mathcal{I}_t = \{I_t^{(m)}\}_{m=1}^M$, and (ii) environmental sensor readings $\mathcal{S}_t = \{s_t^{(n)}\}_{n=1}^N$. Each sensor reading is a seven-channel vector $s_t^{(n)} = [\text{smoke}, \text{temp.}, \text{flame}, \text{water}, \text{gas}, \text{vibration}, \text{tilt}]_t^{(n)}$ (1)

and the synchronized environment-centric observation is

$$z_t = \{\mathcal{I}_t, \mathcal{S}_t\} \quad (2)$$

b) *Robot state (no onboard hazard perception)*: The robot maintains an onboard estimate of its pose x_t for localization and motion execution. Crucially, *no onboard sensing is used for hazard inference*; all reasoning about environmental risk is derived exclusively from z_t . This separation allows hazards to be detected and reasoned about before the robot physically encounters them. Robot state uncertainty and dynamic obstacle interactions are handled by the underlying navigation stack and are not explicitly modeled in the risk belief.

c) *Latent environmental risk state*: We model hazardous operating conditions as a latent, spatio-temporal *environmental risk field*. Let $r_t(\xi) \in \{0, 1\}$ denote whether a location $\xi \in \Omega$ is unsafe at time t . The corresponding belief is defined as

$$B_t(\xi) = p(r_t(\xi) = 1 \mid \cdot) \quad (3)$$

where $B_t(\xi) \in [0, 1]$ represents the probability of being unsafe. Risk evolution is assumed to follow a first-order Markov process

$$p(r_t \mid r_{t-1}) \quad (4)$$

capturing hazard persistence and gradual spatial evolution (e.g., spreading water, sustained heat, or prolonged vibration).

d) *Instantaneous risk observation*: An environment-centric perception model produces an *instantaneous* probabilistic risk observation over the workspace,

$$\hat{Y}_t(\xi) = p(r_t(\xi) = 1 \mid z_t) \quad (5)$$

which is treated as a noisy measurement of the latent risk state. These observations may be incomplete or unreliable due to partial sensor coverage and environmental disturbances.

e) *Latent risk belief*: To avoid reacting to transient fluctuations, we maintain a temporally filtered belief over the latent risk state,

$$b_t(\xi) p(r_t(\xi) = 1 \mid \hat{Y}_{1:t}) \quad (6)$$

yielding a spatial risk belief map $B_t(\xi) = b_t(\xi)$. The specific filtering mechanism used to compute this belief is detailed in Sec. IV.

f) *Risk-aware planning objective*: The filtered risk belief is converted into a planning-compatible cost

$$C_r(\xi) = -\log(1 - B_t(\xi) + \epsilon) \quad (7)$$

where ϵ ensures numerical stability. Given a standard navigation cost $C_{\text{map}}(\pi)$, the planning objective is

$$\pi^* = \arg \min_{\pi} \left(C_{\text{map}}(\pi) + \lambda C_r(\pi) \right) \quad (8)$$

where λ controls the trade-off between efficiency and safety. This formulation penalizes paths that traverse regions with sustained high-risk belief while remaining compatible with standard navigation planners. The resulting path is executed using onboard localization and control, while all hazard inference remains environment-centric.

IV. METHODOLOGY

This section presents the EnvGuard framework, detailing how environment-mounted sensing is used to estimate instantaneous risk, maintain a temporally filtered latent belief, and integrate this belief into standard navigation planning.

A. System Overview and Key Contribution

We propose *EnvGuard*, an environment-centric risk-aware navigation framework for mobile robots operating in hazardous environments. In contrast to conventional robot-centric approaches, hazard perception is performed exclusively by infrastructure-mounted cameras and environmental sensors, while the robot relies on onboard sensing only for localization and motion execution.

Key methodological contribution. The primary contribution of this work is a principled formulation of *environmental navigation risk as a latent, temporally evolving belief* inferred from environment-centric multimodal sensing and integrated into robot planning. Rather than treating risk as an instantaneous score or a static cost derived directly from raw observations, EnvGuard models risk as a hidden environmental state that must be inferred over time under partial observability.

This formulation directly addresses limitations of reactive and instantaneous risk-aware navigation. In real hazardous environments, sensor observations are noisy, spatially incomplete, and often affected by transient disturbances (e.g., smoke plumes or brief temperature spikes). Directly coupling such observations to planning can lead to unstable behavior, excessive replanning, or overly conservative avoidance. EnvGuard therefore decouples *risk inference* from *risk utilization*: instantaneous multimodal predictions are treated as noisy evidence, while navigation decisions are driven by a temporally filtered latent belief that explicitly models hazard persistence and uncertainty.

Crucially, the proposed belief is *environment-centric*. Risk is inferred from infrastructure-mounted sensors that observe regions ahead of the robot, extending the perceptual horizon beyond onboard sensing. The robot itself performs no hazard perception and relies exclusively on the inferred belief for risk-aware decision making. This separation enables proactive navigation decisions before the robot physically encounters unsafe regions.

The overall system architecture, illustrated in Fig. 1, consists of three components: (i) environment-centric multimodal risk perception, (ii) latent risk belief estimation via Bayesian filtering, and (iii) risk-aware planning through costmap integration.

B. Environment-Centric Risk Perception

We adopt the environment-centric sensing formulation defined in Sec. III. At each time step t , synchronized environment-mounted camera images $\mathcal{I}_t$ and environmental sensor readings $\mathcal{S}_t$ are processed to estimate instantaneous navigation risk over the workspace Ω .

a) *Spatial risk grid*: The workspace is discretized into a fixed-resolution 2D grid aligned with the navigation map. All risk estimates are expressed in this grid coordinate frame, and each grid cell $\xi \in \Omega$ represents a local region for which risk is inferred.

b) *Camera-to-grid projection*: Environment-mounted cameras are geometrically calibrated with respect to the workspace. Visual features extracted from each camera view are projected onto the ground plane using known extrinsics [9] (and depth where available), allowing visual evidence to be associated with corresponding grid cells. When multiple cameras observe the same cell, projected features are aggregated using a learned fusion operator.

c) *Sensor preprocessing and grid association*: Raw environmental sensor streams are filtered to suppress high-frequency noise and normalized to a common scale. Each sensor node is associated with nearby grid cells based on its physical location, with influence spatially distributed using a fixed-radius kernel. Missing or delayed sensor readings are explicitly masked and propagated as null tokens to prevent spurious inference.

C. Unified Multimodal Risk Encoder

To integrate heterogeneous environment-centric observations into a unified representation, we employ a multimodal risk encoder that produces an instantaneous probabilistic risk estimate over the workspace. The encoder outputs

$$\hat{Y}_t(\xi) = p(r_t(\xi) = 1 \mid z_t), \quad \xi \in \Omega \quad (9)$$

as defined in Sec. III, where $z_t = \{\mathcal{I}_t, \mathcal{S}_t\}$.

a) *Token construction and fusion*: Visual features are extracted from each camera using a ResNet-50 backbone and spatially pooled into visual tokens aligned with the 2D grid. Environmental sensor readings are projected through a multi-layer perceptron to form sensor tokens. Visual and sensor tokens are concatenated into a single token sequence, and learned positional encodings are added to preserve spatial and temporal structure.

b) *Attention-based multimodal encoding*: The unified token sequence is processed by a transformer encoder composed of multi-head self-attention, residual connections, layer normalization, and feed-forward networks. Early cross-modal attention enables interaction between visual and sensor evidence, allowing correlated cues (e.g., heat and visual smoke) to jointly influence risk estimation.

c) *Temporal context modeling*: To capture short-term hazard dynamics while limiting system latency, the encoder processes a sliding temporal window of $K = 5$ recent observations. This window length was chosen empirically to balance responsiveness and stability: shorter windows were overly reactive, while longer windows increased latency without measurable gains. A temporal transformer captures emerging trends across the window.

After the final transformer layer, a token pooling operation aggregates the per-cell token representations, followed by a linear projection to produce the per-cell instantaneous risk estimate $\hat{Y}_t(\xi)$.

d) *Training*: The encoder is trained on labeled hazard data using focal loss to address class imbalance between safe and hazardous regions. The output $\hat{Y}_t(\xi)$ is treated as noisy probabilistic evidence rather than a definitive hazard decision.

D. Latent Risk Belief via Bayesian Filtering

For clarity, $\hat{Y}_t(\xi)$ corresponds to the instantaneous risk estimate shown in Fig. 1, while $b_t(\xi)$ (or equivalently $B_t(\xi)$) denotes the temporally filtered latent risk belief.

We model environmental navigation risk as a latent, time-varying random variable that cannot be directly observed. For each grid cell ξ , the belief is defined as

$$b_t(\xi)p(r_t(\xi) = 1 \mid \hat{Y}_{1:t}) \quad (10)$$

The motivation for Bayesian filtering is that physical hazards such as fire, flooding, or structural instability exhibit strong temporal persistence. Instantaneous predictions may fluctuate due to sensor noise or partial coverage and are therefore unreliable when used directly for planning.

The belief update consists of two steps:

a) *Prediction (State Transition)*: Hazard evolution is modeled as a first-order Markov process:

$$\bar{b}_t(\xi) = p_{11} b_{t-1}(\xi) + p_{01} [1 - b_{t-1}(\xi)] \quad (11)$$

We set $p_{11} = 0.95$ and $p_{01} = 0.05$ to reflect the high persistence of physical hazards once they manifest, while still allowing gradual decay when evidence disappears. These values were selected based on domain knowledge and empirically validated for stability.

b) *Update (Bayesian)*: The predicted belief is updated using the instantaneous risk estimate as soft evidence:

$$b_t(\xi) = \frac{\hat{Y}_t(\xi) \bar{b}_t(\xi)}{\hat{Y}_t(\xi) \bar{b}_t(\xi) + (1 - \hat{Y}_t(\xi)) (1 - \bar{b}_t(\xi))} \quad (12)$$

When a grid cell is unobserved, an uninformative likelihood is applied so that belief evolution is governed primarily by the prediction model.

c) *Planning-compatible belief*: The resulting belief map $B_t(\xi)$ is spatially smooth, bounded, and temporally stable. This produces a planning-compatible risk signal and cost structure that standard grid-based planners (e.g., A*) can efficiently optimize without oscillations or myopic behavior, in contrast to instantaneous or thresholded risk maps.

E. Risk-Aware Planning and Execution

The latent risk belief is converted into a risk costmap using a negative log-likelihood formulation:

$$C_r(\xi) = -\log(1 - b_t(\xi) + \epsilon) \quad (13)$$

where ϵ ensures numerical stability. This transformation penalizes paths that traverse regions with high inferred risk and corresponds to the cost construction illustrated in Fig. 1.

The risk cost is combined with the static navigation costmap:

$$C_{\text{total}}(\xi) = C_{\text{map}}(\xi) + \lambda C_r(\xi) \quad (14)$$

We assume a static baseline navigation costmap C_{map} ; dynamic obstacles are not considered in this work. The safety weight λ is treated as a hyper-parameter and tuned on a held-out validation set to balance mission efficiency and hazard exposure.

a) *Replanning trigger*: Replanning is triggered when

$$\max_{\xi \in \pi_{\text{current}}} |b_t(\xi) - b_{t-\Delta t}(\xi)| > \tau \quad (15)$$

where the threshold τ is selected empirically to ensure timely response without excessive replanning.

The rover executes the planned path using onboard localization (e.g., IMU and pose estimation) and motion control. No onboard hazard perception or dynamic obstacle reasoning is performed.

Practical assumptions. EnvGuard assumes sufficient infrastructure-sensor coverage, reliable time synchronization, and calibration between environment-mounted sensors and the navigation map. These assumptions enable proactive risk estimation beyond the robot’s onboard perceptual range, but performance may degrade in unobserved regions, under communication delay, or when camera/sensor calibration is inaccurate.

b) *Summary*: By explicitly modeling environmental risk as a latent, temporally filtered belief, EnvGuard converts noisy environment-centric sensing into a stable, planning-compatible signal. This enables proactive and reliable navigation in hazardous environments using standard planning algorithms without modification, paving the way for the experimental validation presented in the next section.

V. EXPERIMENTS AND RESULTS

This section evaluates the proposed environment-centric risk-aware navigation framework on a *real physical rover operating in a multi-hazard testbed*. The goal is to assess whether modeling environmental risk as a *latent, temporally filtered state* enables safer and more stable navigation compared to risk-unaware, reactive, and existing risk-aware alternatives. All hazard perception is performed exclusively using environment-mounted cameras and sensors, while the rover relies on onboard sensing only for localization and motion execution. All results are obtained from real sensor logs and executed rover trajectories.

A. Physical Multi-Hazard Testbed

We constructed a controlled indoor physical multi-hazard testbed to evaluate navigation under partial observability and dynamically evolving environmental risks. The testbed comprises modular corridor panels forming three geometrically similar routes (Paths A, B, and C) between fixed start and goal locations (Fig. 2), ensuring that performance differences arise from hazard conditions rather than structural bias. Each corridor can be independently instrumented with hazards such as water accumulation, localized thermal sources, vibration, and unstable ground, which are activated dynamically during execution. Hazardous regions are positioned ahead of the rover and are not directly observable via onboard sensing prior to

TABLE I: Environment-mounted sensing infrastructure in the physical testbed

Sensor	Modality	Accuracy/Res.	Placement
MQ-2	Smoke	$\pm 10\%$	Hazard enclosure
DS18B20	Temperature	$\pm 0.5^\circ\text{C}$	Hazard enclosure
KY-026	Flame	Binary	Hazard enclosure
HW-038	Water	Binary	Floor-mounted
MQ-135	Gas	$\pm 15\%$	Hazard enclosure
801S	Vibration	Binary	Floor-mounted
MPU6050	Tilt	$\pm 0.1^\circ$	Floor-mounted
OAK-D Pro	RGB-D	640×480	Corridor corners




Fig. 2: Robot traversal sequences showing three key phases: (a-c) First traversal set and (d-f) second traversal set. The robot navigates from start through mid-sections to exit points in hazardous environments.

entry, explicitly testing *anticipatory* navigation enabled by environment-centric perception.

The environment is equipped with fixed RGB-D cameras and distributed sensor nodes measuring smoke, temperature, gas concentration, vibration, and tilt, summarized in Table I. Sensor coverage is partial and overlapping, reflecting realistic infrastructure-assisted deployments. Experiments were conducted over multiple days with repeated rover traversals under diverse hazard schedules, during which synchronized sensor streams, latent risk beliefs, planner outputs, and executed trajectories were recorded. The rover platform used in all experiments was fully designed and fabricated in-house, including a custom 3D-printed chassis and tailored onboard hardware, and operates in both outdoor and indoor environments such as uneven terrain and stair traversal (Fig. 3).

B. Latent Risk Belief Under Real Sensor Noise

We first analyze the behavior of the latent risk belief independent of navigation. Fig. 4a compares instantaneous risk observations produced by the multimodal encoder with the Bayesian-filtered belief under real sensor noise. Instantaneous predictions exhibit high-frequency fluctuations caused by sensor noise, environmental turbulence, and transient disturbances. Direct use of such signals for planning leads to unstable and overly reactive behavior. In contrast, Bayesian



(a) Outdoor uneven terrain (b) Indoor stair traversal

Fig. 3: Rover platform operating in outdoor and indoor environments.

TABLE II: Risk belief stability under real sensor noise

Estimate	MSE ↓	Delay (s) ↓	False Positives ↓
Instantaneous	0.187	0.3	0.21
Filtered (Ours)	0.032	1.1	0.05

filtering produces a smooth belief trajectory that preserves hazard onset and persistence while suppressing spurious spikes.


Quantitatively, the filtered belief reduces mean squared error by more than five times relative to instantaneous predictions, while introducing a modest detection delay of approximately one second. Table II quantitatively evaluates the effect of Bayesian filtering on risk belief stability under real sensor noise using standard accuracy and false-positive metrics. This delay is acceptable at the navigation level and provides a favorable trade-off between responsiveness and stability.

C. Risk Evolution Across Alternative Paths


Fig. 4b shows the evolution of the filtered risk belief aggregated along Paths A, B, and C during a representative physical run. Each curve corresponds to the belief associated with a candidate corridor. Path B, instrumented with a water hazard, exhibits an early rise in risk, followed by Path A with a thermal hazard. Path C remains largely safe until late activation. Warning ($\tau = 0.33$) and unsafe ($\tau = 0.7$) thresholds are shown explicitly and provide actionable signals for monitoring and replanning. Across repeated runs, the warning threshold is crossed 4.2 ± 1.3 s before physical hazard contact, providing sufficient temporal margin for proactive replanning decisions.

D. Online Path Switching in the Physical Testbed

Fig. 5a illustrates the risk experienced by the rover during execution, with background shading indicating the selected corridor. Vertical dashed markers denote replanning events triggered by changes in the latent risk belief. The rover initially commits to Path B, switches to Path C when the risk exceeds the warning threshold, and replans again as hazards propagate. All switching decisions occur before the rover enters unsafe regions, confirming that the behavior is proactive rather than reactive. Fig. 5b summarizes the system-level decision timeline, including hazard activations, belief threshold crossings, replanning events, and mission completion. This visualization



(a) Bayesian filtering of noisy risk observations



(b) Risk belief evolution across candidate paths

Fig. 4: Latent risk belief estimation using environment-centric sensing. (a) Bayesian filtering suppresses noisy instantaneous risk observations and tracks the latent hazard state under 5% sensor noise. (b) Filtered risk belief evolution across candidate paths; warning ($\tau = 1/3$) and unsafe ($\tau = 0.7$) thresholds trigger monitoring, replanning, and emergency avoidance.

TABLE III: End-to-end system latency on the physical platform

Component	Mean (ms)	Std (ms)
Multimodal Encoding	65	8
Bayesian Filtering	8	2
Risk Costmap Update	5	1
Path Planning	22	7
Total	95	15


highlights the causal relationship between environment-centric risk perception, Bayesian belief estimation, and navigation decisions in the physical testbed.

E. System Latency and Real-Time Performance


Table III reports end-to-end latency measured on the physical system. The complete pipeline operates at 10Hz with an average latency of 95 ms, demonstrating feasibility for real-time deployment.

F. Quantitative Comparison with Baselines

We compare four navigation strategies: risk-unaware (*NoRisk*), threshold-based reactive (*Reactive*), instantaneous



(a) Risk-driven online path switching during navigation



(b) Decision timeline and replanning events

Fig. 5: Risk-aware navigation behavior enabled by latent belief estimation. (a) Rover traversal with online path switching driven by the filtered risk belief. Background shading indicates the active corridor, while vertical markers denote replanning events triggered before entering unsafe regions. (b) System-level decision timeline illustrating hazard activations, belief threshold crossings, replanning actions, and mission completion.

TABLE IV: Navigation performance in the physical testbed

Method	Success $\uparrow$	Exposure (s) $\downarrow$	Replans $\downarrow$
NoRisk (No hazard modeling)	0.82	42.5	0.4
Reactive (Threshold-based [13, 3])	0.88	28.3	4.7
Probabilistic (Instantaneous risk [14, 2])	0.90	19.8	3.9
EnvGuard (Ours)	0.95	12.8	2.1

probabilistic (*Probabilistic*), and the proposed method (*EnvGuard*). The NoRisk baseline ignores hazard information entirely and plans solely with geometric cost. Reactive uses instantaneous risk estimates with threshold-based replanning. Probabilistic incorporates unfiltered risk probabilities directly as a planning cost without temporal belief filtering. EnvGuard uses environment-centric sensing with Bayesian risk belief estimation. EnvGuard achieves the highest success rate and lowest hazard exposure, while avoiding the excessive replanning observed in reactive approaches. Table IV reports quantitative navigation performance measured on the real physical testbed, comparing EnvGuard against risk-unaware, reactive, and probabilistic baselines. Fig. 6 provides a detailed

TABLE V: Comparison with representative risk-aware navigation methods

Method	Risk Source	Temporal	Proactive	Physical
RAMP [13]	Robot	No	Partial	No
EVORA [3]	Robot	No	No	No
Probabilistic Planning [2]	Robot	Probabilistic	Partial	No
EnvGuard (Ours)	Environment	Bayesian	Yes	Yes

quantitative comparison between EnvGuard and competing navigation strategies across complementary performance metrics. Fig. 6a summarizes overall system behavior using a multi-criteria radar plot, highlighting that EnvGuard achieves a more balanced trade-off between safety, efficiency, and replanning stability than reactive and probabilistic baselines.

Fig. 6b compares hazard exposure time across methods. EnvGuard significantly reduces the duration spent in unsafe regions, confirming that temporally filtered, environment-centric risk beliefs enable earlier and more consistent avoidance of hazards. Fig. 6c reports mission success rates. While risk-unaware navigation frequently fails due to hazardous encounters, EnvGuard maintains the highest success rate, demonstrating that safety gains do not come at the expense of task completion. Finally, Fig. 6d presents statistical significance results from paired tests across physical runs. The improvements achieved by EnvGuard over all baselines are statistically significant, indicating that the observed gains are robust and not attributable to random variation.

G. Comparison with representative risk-aware navigation systems

Table V positions EnvGuard relative to representative risk-aware navigation systems. Prior methods such as RAMP [13] and EVORA [3] estimate risk primarily from *robot-centric perception* and treat risk either as an instantaneous signal or as a learned traversability score, without explicit temporal belief maintenance. Probabilistic planning approaches [2] incorporate uncertainty during planning, but typically assume risk models derived from onboard sensing and do not explicitly maintain a persistent environmental risk belief across time.

In contrast, EnvGuard uniquely combines (i) *environment-centric risk inference* using infrastructure-mounted sensing, (ii) *Bayesian temporal filtering* of latent environmental risk, and (iii) *proactive navigation decisions* validated in a real physical multi-hazard testbed. This combination enables earlier hazard anticipation and more stable replanning behavior, which is not supported by existing methods.

Importantly, EnvGuard is the only approach among those compared that demonstrates proactive risk-aware navigation using environment-level sensing with physical experimental validation.

H. Summary of Experimental Findings

Physical experiments demonstrate that Bayesian filtering produces stable latent risk beliefs under real sensor noise, environment-centric sensing enables proactive avoidance of hazards before physical exposure, and the proposed framework




Fig. 6: Quantitative comparison with state-of-the-art methods. EnvGuard consistently achieves lower hazard exposure, higher mission success, and statistically significant improvements over risk-unaware, reactive, and probabilistic baselines.

consistently outperforms risk-unaware, reactive, and probabilistic baselines. These results validate the central claim of this work: *treating environmental risk as a latent, temporally filtered state enables safer robot navigation without onboard hazard perception*.

VI. CONCLUSION

This paper introduced *EnvGuard*, an environment-centric risk-aware navigation framework that enables proactive hazard avoidance by decoupling hazard perception from the robot. Environmental risk is modeled as a latent state inferred from infrastructure-mounted sensing and stabilized via Bayesian filtering into a planning-compatible belief. Experiments on a physical rover in a multi-hazard testbed demonstrate that the latent risk belief suppresses noise while preserving hazard onset, allowing danger to be anticipated seconds before exposure. EnvGuard significantly reduces hazard exposure compared to risk-unaware, reactive, and instantaneous risk-based baselines, while maintaining high task success.

Two key advantages distinguish EnvGuard: (i) risk is inferred from environment-mounted sensors, extending perception beyond the robot, and (ii) risk is treated as a temporally evolving belief rather than a reactive signal. This enables earlier, more stable avoidance in dynamic settings. Overall, environment-centric perception with principled temporal inference provides a practical pathway toward safer autonomous navigation in hazardous environments.

Limitations and future directions. While EnvGuard supports online replanning in response to dynamically evolving environmental hazards, the current framework assumes calibrated infrastructure sensors with sufficient spatial coverage and reliable communication with the robot. These assumptions enable proactive risk estimation beyond the robot’s onboard sensing range, but performance may degrade in unobserved areas, under communication delay, or when camera/sensor calibration is inaccurate. The current framework focuses on environment-level risk and does not explicitly model fast-moving obstacles such as people or other robots. In addition, the Bayesian transition parameters governing hazard persistence are fixed and selected based on domain knowledge rather than learned online. Future work could extend the framework to jointly reason about environment-centric risk and dynamic obstacle avoidance, incorporate adaptive or learned hazard dynamics, and support coordinated multi-robot navigation under shared risk beliefs.

ACKNOWLEDGMENT

This work was supported in part by the Institute of Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korea Government (MSIT, Development of Artificial Intelligence Based Video Security Technology and Systems for Public Infrastructure Safety) under Grant RS-2019-II190231, and in part by the National Research Foundation of Korea (NRF) grant funded by the Korea Government (MSIT) under Grant RS-2024-00335672.

REFERENCES

- [1] Aaron D. Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control Barrier Functions: Theory and Applications. *arXiv preprint*, 2019. doi: 10.48550/arXiv.1903.11199. URL <https://arxiv.org/abs/1903.11199>.
- [2] Lars Blackmore, Masahiro Ono, and Brian C. Williams. Chance-Constrained Optimal Path Planning with Obstacles. *IEEE Transactions on Robotics*, 27(6):1080–1094, 2011. doi: 10.1109/TRO.2011.2161160. URL <https://doi.org/10.1109/TRO.2011.2161160>.
- [3] Xiaoyi Cai, Siddharth Ancha, Lakshay Sharma, Michael Everett, and Jonathan P. How. EVORA: Deep Evidential Traversability Learning for Risk-Aware Off-Road Autonomy. *IEEE Transactions on Robotics*, 40:3756–3777, 2024. doi: 10.1109/TRO.2024.3431828. URL <https://doi.org/10.1109/TRO.2024.3431828>.
- [4] Peter Corke, Tim Wark, Raja Jurdak, Wen Hu, Philip Valencia, and Darren Moore. Environmental Wireless Sensor Networks. *Proceedings of the IEEE*, 98(11):1903–1917, 2010. doi: 10.1109/JPROC.2010.2068530. URL <https://doi.org/10.1109/JPROC.2010.2068530>.
- [5] Peter Corke, Wen Hu, and Matthew Dunbabin. An RPC-Based Service Framework for Robot and Sensor Network Integration. In *Proceedings of the IEEE Vehicular Technology Conference (VTC Spring)*, pages 1–6, 2011. doi: 10.1109/VETECS.2011.5956411. URL <https://doi.org/10.1109/VETECS.2011.5956411>.
- [6] Alberto Elfes. Using Occupancy Grids for Mobile Robot Perception and Navigation. *Computer*, 22(6):46–57, 1989. doi: 10.1109/2.30720. URL <https://doi.org/10.1109/2.30720>.
- [7] Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. doi: 10.1109/TSSC.1968.300136. URL <https://doi.org/10.1109/TSSC.1968.300136>.
- [8] Hailong Huang, Andrey V. Savkin, Ming Ding, and Chao Huang. Mobile Robots in Wireless Sensor Networks: A Survey on Tasks. *Computer Networks*, 148:1–19, 2019. doi: 10.1016/j.comnet.2018.10.018. URL <https://doi.org/10.1016/j.comnet.2018.10.018>.
- [9] Palash Y. Ingle and Young-Gab Kim. Multi-Sensor Data Fusion Across Dimensions: A Novel Approach to Synopsis Generation Using Sensory Data. *Journal of Industrial Information Integration*, page 100876, 2025. doi: 10.1016/j.jii.2025.100876. URL <https://doi.org/10.1016/j.jii.2025.100876>.
- [10] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence*, 101(1–2):99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X. URL <https://www.sciencedirect.com/science/article/pii/S000437029800023X>.
- [11] Ben Kehoe, Sachin Patil, Pieter Abbeel, and Ken Goldberg. A Survey of Research on Cloud Robotics and Automation. *IEEE Transactions on Automation Science and Engineering*, 12(2):398–409, 2015. doi: 10.1109/TASE.2014.2376492. URL <https://ieeexplore.ieee.org/document/7006734>.
- [12] Robin R. Murphy. *Disaster Robotics*. The MIT Press, 2014. doi: 10.7551/mitpress/9407.001.0001. URL <https://direct.mit.edu/books/monograph/3408/Disaster-Robotics>.
- [13] Lakshay Sharma, Michael Everett, Donggun Lee, Xiaoyi Cai, Philip Osteen, and Jonathan P. How. RAMP: A Risk-Aware Mapping and Planning Pipeline for Fast Off-Road Ground Robot Navigation, 2023. URL <https://arxiv.org/abs/2210.06605>. arXiv:2210.06605.
- [14] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. The MIT Press, 2005. URL <https://mitpress.mit.edu/9780262303804/probabilistic-robotics/>.