computation

MDPI

Article

Physiological Signal-Guided Uncertainty Management for Autonomous UAVs in Human–UAV Supervisory Control

Jun Che [1,2], Feng Zhu [1] and Hanbin Xiao [1,*]

1 School of Transportation and Logistics Engineering, Wuhan University of Technology, Wuhan 430063, China; 307595@whut.edu.cn (J.C.); zhufeng@whut.edu.cn (F.Z.)
2 School of Logistics, Wuhan Technology and Business University, Wuhan 430065, China
* Correspondence: xhbchina@126.com

Abstract

This paper presents a physiological signal-guided uncertainty-management framework that integrates real-time indicators of the operator's supervisory state into UAV autonomy to improve obstacle avoidance and risk-aware decision-making under uncertain conditions. During UAV supervisory control, multimodal physiological signals, including heart rate variability, blood pressure, electrodermal activity, and other cardiovascular or stress-related measures, are time-synchronized with UAV telemetry, perceived obstacle fields, planner confidence, environmental uncertainty estimates, and operator intervention logs, including waypoint edits, overrides, and replanning commands. These heterogeneous data streams are fused using a Bayesian hierarchical state-space framework to estimate latent supervisory states representing trust miscalibration, risk sensitivity, and situational-awareness degradation. The estimated states are then incorporated into the UAV decision-making stack as bounded uncertainty-management parameters that regulate safety margins, replanning priority, and risk preference without relaxing hard safety constraints. The framework was evaluated in a simulation-based human-in-the-loop study involving 24 operators and 180 UAV obstacle avoidance missions. Performance was assessed using held-out-operator mission success AUROC, proxy-state RMSE, mission success rate, mean minimum obstacle clearance, mission risk index, operator override rate, and replanning latency. The results support the feasibility of using physiological, behavioral, vehicle, and environmental information to adapt UAV supervisory control to uncertainty in both the operating environment and human supervisory readiness.

Keywords: human–UAV interaction; obstacle avoidance; physiological data; machine learning

Academic Editor: Garzia Fabio

Received: 7 June 2026
Revised: 12 August 2026
Accepted: 17 August 2026
Published: 20 August 2026

1. Introduction

Autonomous unmanned aerial vehicles (UAVs) are increasingly deployed for last-mile and middle-mile logistics, operating in cluttered, dynamic low-altitude environments populated by buildings, vegetation, power lines, and other airspace users. Robust obstacle avoidance under uncertainty is central to the safety and scalability of these operations. Yet, despite substantial advances in autonomous planning, current methods remain fundamentally limited in how they handle the human element, and this gap is emerging as a critical barrier to real-world deployment [1–4].

State-of-the-art UAV autonomy integrates onboard perception (depth cameras, LiDAR, radar), state estimation (GPS/INS fusion, visual-inertial odometry), and real-time planning algorithms—including velocity obstacles, model predictive control (MPC), and sampling-based methods such as RRT*—to generate collision-free trajectories. Risk-aware extensions such as chance-constrained MPC, robust optimization, and Monte Carlo tree search have further improved the system's ability to reason about sensing and modeling uncertainty. These are meaningful advances, but they share a critical architectural assumption: the human supervisor is treated as an external safety layer rather than as an integral source of contextual information. This assumption creates several compounding problems. First, dynamic environments introduce rapid, unpredictable changes—sudden wind gusts, unexpected aerial traffic, or moving ground obstacles that exceed the reactive capacity of model-based planners calibrated on static or slow-changing risk assumptions [5,6]. Planners that rely on predefined uncertainty budgets cannot adapt when environmental volatility spikes mid-mission. Second, current obstacle avoidance stacks apply fixed safety margins that are either too conservative (reducing route efficiency and throughput) or insufficiently conservative (increasing collision risk) because they have no mechanism for sensing the actual level of situational awareness on the human side. Third, and most critically, these systems cannot distinguish between a calm, attentive supervisor and an overloaded operator at risk of missing critical cues even when the observable command patterns of both operators appear similar. This mismatch leads to miscalibrated trust in automation, delayed or absent interventions during rapidly evolving obstacle events, and poor scalability as operators are tasked with supervising larger fleets [7,8].

Existing human-in-the-loop architectures partially address these limitations by enabling operators to monitor multiple vehicles, approve high-level plans, and intervene during off-nominal events. Shared authority mechanisms—including mixed-initiative waypoint editing, priority zones, no-fly volumes, and modal autonomy switching—have extended operator influence over trajectory generation [9,10]. However, these interfaces treat human input as discrete commands or explicit mode switches. They provide no continuous model of the operator's cognitive or affective state. As a result, the system's obstacle avoidance behavior remains static with respect to human factors: it does not know whether the operator is attentive, stressed, fatigued, or cognitively overloaded. When a drone operator misses a critical visual prompt due to post-task fatigue, or fails to process high-dimensional sensor data quickly enough during a time-critical delivery, the autonomy stack has no mechanism to compensate [11]. Compounding this is the difficulty of conveying operator intent under uncertainty. Users often struggle to express nuanced preferences—for example, how to balance safety buffers against route efficiency, or whether to avoid a particular corridor for reasons not captured in the map. Multimodal interaction remains underdeveloped for UAV logistics contexts, limiting the operator's ability to communicate risk judgments in real time [12,13].

A growing body of research in human–automation interaction has demonstrated that physiological signals provide reliable, continuous proxies for mental workload, vigilance, stress, and trust. Metrics derived from heart rate variability (HRV) [14–16], electrodermal activity, blood pressure, and eye behavior (pupil dilation, fixation patterns) correlate robustly with cognitive load, risk sensitivity, and decision-making capacity under uncertainty. Adaptive automation frameworks in aviation and ground vehicle domains have used such signals to modulate task allocation and interface complexity—offloading tasks when workload exceeds thresholds, or simplifying displays when overload is detected. Yet these mechanisms have rarely been integrated with the core motion planning and obstacle avoidance logic of autonomous UAVs. The signals are monitored but not acted upon where it matters most: inside the planner.

This gap is precisely what the proposed framework closes. Rather than treating human physiological state as a monitoring curiosity, the approach embeds it as a feedback signal that directly modulates obstacle avoidance behavior in real time. When an operator is physiologically stable and demonstrably attentive, the system can safely relax conservative safety margins to recover route efficiency. When biometric indicators signal elevated stress, reduced vigilance, or degraded situational awareness, the obstacle avoidance module responds autonomously, inflating safety buffers, biasing trajectory sampling toward lower-risk paths, tightening MPC separation constraints, and reducing speed in high-clutter regions—without requiring an explicit operator command that the overloaded operator may be unable to issue in time. This is effective for several mutually reinforcing reasons. First, it makes the obstacle avoidance stack reactive to the real uncertainty in the system, which includes not only environmental dynamics but also the reliability of human oversight at any given moment. Second, it removes the dependency on accurate operator self-reporting or explicit mode switching, both of which degrade under stress and workload. Third, it enables autonomy to compensate precisely when human performance is most likely to fail during high-clutter, time-critical segments without permanently removing the operator from meaningful control. Fourth, by fusing biometric features with contextual variables (obstacle density, sensor noise, fleet size) in hierarchical Bayesian models, the system can infer latent cognitive states—risk sensitivity, trust in autonomy and situational awareness—that evolve across missions and inform adaptive policies with increasing accuracy over time [17–19].

This work proposes a physiology-aware uncertainty-management framework (Figure 1) that directly links operator cognitive state, inferred from multimodal biometric signals to the real-time obstacle avoidance behavior of autonomous logistics UAVs. In the envisioned system, an operator supervises one or more UAVs through a mixed-initiative interface in which control authority is dynamically partitioned between human and autonomy. Physiological signals including time-domain and frequency-domain features of heart rate variability (HRV) such as RMSSD, SDNN, and the LF/HF spectral ratio, as well as continuous blood pressure waveforms, electrodermal activity (EDA), and eye-tracking metrics including pupil diameter and fixation entropy are acquired at high sampling rates ($\geq$250 Hz for cardiovascular signals) via wearable sensors and processed through artifact rejection and bandpass filtering pipelines. These physiological streams are then continuously time-aligned, using hardware-synchronized timestamps, with UAV telemetry (UAV pose estimates, perceived obstacle positions and velocities, onboard sensor confidence scores, and planner replanning frequency) and human decision logs (waypoint insertion and deletion events, override commands, authority mode switches, and reaction latencies) [20–22].

The proposed architecture integrates human-state sensing, UAV perception, multimodal data fusion, and adaptive flight control within a closed supervisory loop. Layer 1 acquires quality-controlled physiological features and records waypoint edits, route changes, altitude commands, overrides, replanning requests, and confidence ratings. Layer 2 provides UAV pose, velocity, attitude, obstacle proximity and density, local map uncertainty, sensor quality, and planner confidence. Layer 3 aligns these streams using timestamp and synchronization-quality metadata and estimates a posterior distribution over operator workload, trust calibration, risk sensitivity, and situational awareness together with autonomy reliability and mission risk. Layer 4 maps the fused state to bounded adaptation commands for autonomy level, replanning priority, risk preference, speed scaling, and controller tuning. These commands can make the trajectory planner and model predictive controller more conservative, but they cannot weaken hard collision avoidance, separation, actuator, or vehicle-dynamics constraints. Executed motion and updated observations close the loop, and inadequate physiological signal quality triggers the validated physiology-independent fallback controller.

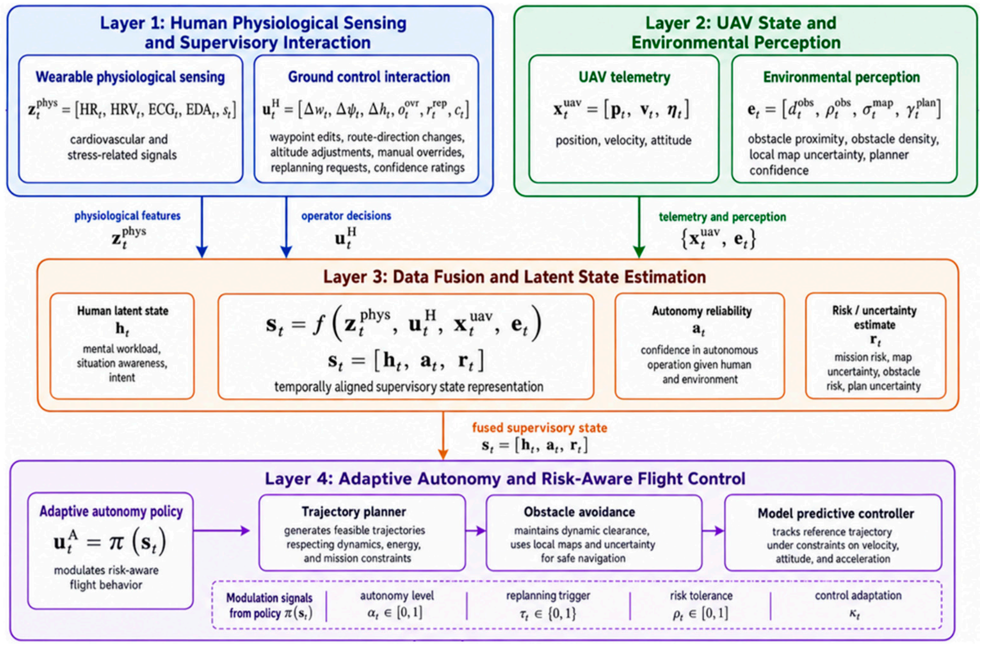

Figure 1. Proposed system architect for physiological signal-guided human–UAV supervisory control.

Feature extraction operates on sliding temporal windows (e.g., 30–60 s epochs with 50% overlap) to produce low-latency biometric feature vectors that capture both tonic baseline shifts and phasic responses to discrete obstacle events. These feature vectors are fused with mission-contextual variables including local obstacle density (derived from occupancy grid entropy), instantaneous sensor noise estimates (e.g., LiDAR return confidence, GPS dilution of precision), number of simultaneously supervised vehicles, and time on task within a hierarchical Bayesian state-space model. The hierarchical structure allows the model to share statistical strength across operators and missions at the population level while simultaneously maintaining individualized posterior distributions over latent cognitive states for each operator. Specifically, the model jointly infers three latent constructs that evolve as continuous hidden Markov processes: (1) cognitive load, indexed primarily by HRV spectral features and pupil dilation; (2) risk sensitivity, estimated from the covariance between physiological arousal indicators and operator override frequency in obstacle-dense segments; and (3) situational-awareness degradation, inferred from fixation entropy, reaction latency distributions, and deviations from expected decision patterns given current environmental complexity. The posterior estimates of these latent states are updated recursively using a particle filter or variational inference scheme to maintain real-time tractability, and are then passed as adaptive parameters into the UAV autonomy stack. Concretely, the latent state estimates modulate four interconnected components of the obstacle avoidance pipeline: (i) the safety buffer inflation factor applied to obstacle bounding volumes in the signed distance field used by the trajectory optimizer; (ii) the risk bias term in the sampling distribution of the RRT*- or MPPI-based motion planner, which shifts trajectory proposals toward lower-risk homotopy classes; (iii) the constraint tightening coefficient on minimum separation distance within the MPC receding-horizon formulation; and (iv) a speed scaling factor that reduces nominal cruise velocity in high-clutter regions when operator state indicators fall below confidence thresholds. These four parameters are jointly governed by a continuous-valued policy function trained offline via Bayesian optimization over historical mission data that maps the inferred latent state vector to a Pareto-optimal operating point on the safety-efficiency frontier given current environmental conditions.

The anticipated benefits are threefold: enhanced safety through real-time coupling of obstacle avoidance aggressiveness to operator resilience indicators; improved mission efficiency through adaptive balancing of safety margins and route optimality rather than static worst-case design; and scalable human supervision by enabling autonomy to compensate for operator overload while generating richer data to inform interface design, training, and workload allocation [23]. These benefits arise not from any single component in isolation, but from the coupling architecture itself: while the multimodal physiological sensing pipeline draws on established biometric processing techniques, and the Bayesian hierarchical state-space model provides a principled formulation for inferring latent operator cognitive states from these signals, the core novelty of this work lies in the adaptive autonomy policy that translates these inferred states into obstacle avoidance behavior, and in the tightly integrated, time-synchronized architecture that closes the loop between operator physiology and UAV planning in real time. It is this end-to-end integration—rather than the sensing or inference components individually—that provides a principled, data-driven pathway for embedding biometric information into the core obstacle avoidance logic of autonomous logistics UAVs, advancing both safety and operational performance in real-world deployment.

2. Materials and Methods

2.1. System Architecture for Physiological Signal-Guided Human–UAV Supervisory Control

The experimental system is designed to investigate how physiological indicators of operator cognitive state can be incorporated into autonomous UAV uncertainty management during human–UAV supervisory control. The architecture integrates multimodal human sensing, operator intervention logging, UAV telemetry, and environmental perception within a time-synchronized mixed-initiative control framework, in which the human operator supervises one or more multirotor logistics UAVs operating in obstacle-rich environments while the autonomy module performs trajectory planning, obstacle avoidance, and replanning under uncertainty. The central objective is to determine whether physiological signals including heart rate, heart rate variability, blood pressure, and operator-reported confidence can provide real-time estimates of supervisory readiness, stress, workload, and risk sensitivity sufficient to adapt UAV safety margins and planning behavior in a principled manner [24–26].

The system architecture consists of four coupled functional layers. The first is the human physiological sensing and supervisory interaction layer, in which wearable sensors continuously acquire cardiovascular and stress-related signals while the ground control interface records all explicit operator inputs, including waypoint modifications, route-direction changes, altitude adjustments, manual overrides, replanning requests, and confidence ratings associated with perceived autonomy reliability. The second is the UAV state and environmental perception layer, which collects onboard telemetry and sensor-derived environmental information, including UAV position, velocity, attitude, obstacle proximity, obstacle density, local map uncertainty, and planner confidence. The third is the data fusion and latent state estimation layer, in which physiological features, decision events, and UAV environment variables are temporally aligned and encoded into a unified supervisory state representation. The fourth is the adaptive autonomy layer, where the inferred human-state and environmental uncertainty variables are passed to the trajectory planner, obstacle avoidance module, and model predictive controller to modulate risk-aware flight behavior [27–29].

All data streams are implemented through a publish–subscribe middleware (ROS2) to support modular integration between physiological sensing, UAV control, and data logging. A common timing framework based on high-resolution timestamps and network-level clock synchronization aligns biometric signals, operator actions, UAV telemetry, and perception outputs. This synchronization is essential because the proposed model seeks to capture the precise temporal relationship between operator physiological response, perceived environmental uncertainty, human intervention behavior, and subsequent UAV planning outcomes.

2.2. Human Physiological Sensing and Supervisory Decision Acquisition

During each experimental trial, the operator supervises UAV navigation through environments containing both scripted and stochastic obstacle configurations. Physiological data are collected continuously using commercial wearable and non-invasive monitoring devices. Heart rate and interbeat interval (IBI) signals are obtained from a chest strap or wrist-worn sensor [30–32] sampled at 100–250 Hz. Blood pressure is recorded using either a continuous non-invasive blood pressure monitor or an intermittent cuff-based device, providing systolic blood pressure (SBP), diastolic blood pressure (DBP), and mean arterial pressure (MAP) [33,34]. Continuous blood pressure measurements are sampled at approximately 0.5–1 Hz, while intermittent measurements are collected every 30–60 s, depending on device capability and experimental constraints.

Physiological and device-quality streams are published as standardized ROS 2 topics for synchronized logging. Operator confidence is collected at predefined mission phases on a 7-point scale [35,36] and is combined with behavioral features such as response latency, recent overrides, and acceptance of suggested trajectories [37–41]. Sparse ratings are assimilated as observations whose uncertainty grows with time since the last rating; they are not spline-interpolated. Between ratings, a latent confidence estimate is updated from behavioral features and a quality-controlled arousal index using the logistic-normal observation model in Equation (1).

$$\text{logit}(\hat{c}_t) = \mathbf{w}^{\mathrm{T}}\mathbf{f}_t + \beta a_t + \epsilon_t,\ \ \epsilon_t \sim \mathcal{N}(0, \sigma_c^2) \tag{1}$$

where the behavioral feature vector and arousal index are evaluated at time t; the regression coefficients are learned from the development cohort; and the Gaussian residual represents uncertainty on the logit scale. The reported confidence estimate is the posterior mean after filtering and is clipped numerically only for display, not by placing noise directly inside a deterministic sigmoid.

2.3. Human Decision-Making and Intervention Logging

Human decisions that modify the UAV's trajectory are treated as discrete, semantically labeled events captured by the ground control interface and encoded using a unified schema [42–44]:

$$e_k = \{t_k, \text{uav\_id}_k, \text{event\_type}_k, \text{payload}_k\} \tag{2}$$

where the timestamp is expressed on the common ground-station time base, the vehicle identifier specifies the supervised UAV, the event type identifies the operator action, and the payload stores action-specific data [45–47]. A waypoint-update payload contains the waypoint identifier, the original coordinates (x_{old}, y_{old}, z_{old}), the revised coordinates (x_{new}, y_{new}, z_{new}), and an optional reason code. The Euclidean displacement in Equation (3) is used as a descriptive intervention-magnitude feature; it is not interpreted as the quality or safety of the correction.

$$\Delta d_k = \parallel \mathbf{p}_k^{\text{new}} - \mathbf{p}_k^{\text{old}} \parallel_2 = \sqrt{\left(x_k^{\text{new}} - x_k^{\text{old}}\right)^2 + \left(y_k^{\text{new}} - y_k^{\text{old}}\right)^2 + \left(z_k^{\text{new}} - z_k^{\text{old}}\right)^2} \tag{3}$$

Path-direction-change events are logged when the operator selects a new corridor, reverses traversal direction, or reorients a path segment. Because headings are circular variables, the change in mean heading is computed with the wrapped angular difference in Equation (4), rather than by unwrapped subtraction.

$$\Delta\psi_k = \mathrm{atan2}\Big(\sin(\bar{\psi}_k^{\mathrm{post}} - \bar{\psi}_k^{\mathrm{pre}}), \cos(\bar{\psi}_k^{\mathrm{post}} - \bar{\psi}_k^{\mathrm{pre}})\Big) \in (-\pi, \pi] \tag{4}$$

where the pre-change and post-change mean headings are in radians. Altitude-change events record the previous commanded altitude and the new commanded altitude; their signed difference is defined by Equation (5).

$$\Delta z_k = z_k^{\mathrm{cmd}} - z_k^{\mathrm{prev}} \tag{5}$$

All intervention events are timestamped, forwarded to the planning service, and stored with the concurrent physiological posterior, signal-quality indicators, planner confidence, and environmental uncertainty. This schema supports model training and real-time adaptation without assigning causal meaning to an isolated physiological response or intervention.

2.4. Multimodal Time-Series Representation and Feature Engineering

Each mission is represented on a common 1 s analysis grid. Continuous measurements are anti-aliased before resampling; persistent state variables use zero-order hold; and discrete operator actions are binned as event indicators, counts, or magnitudes within each interval. Sparse confidence ratings are assimilated by the state estimator with time-varying uncertainty rather than interpolated. NASA-TLX [48–50] is a post-task outcome and is used only for offline validation or covariate analysis, not as an online feature.

The biometric feature vector is constructed from artifact-screened cardiovascular measurements and signal-quality or missingness indicators. Instantaneous heart rate is derived only when the RR interval is expressed in seconds. RMSSD and SDNN are computed over causal 60–120 s windows and interpreted as ultra-short-term features; LF/HF is excluded from the 60 s online vector because reliable spectral estimation generally requires a longer approximately stationary record. Beat-to-beat blood pressure features are included only when a continuous monitor is used. Pulse pressure and mean arterial pressure are defined in Equation (6).

$$\begin{gathered} \mathbf{b}_t = [HR_t, RR_t, RMSSD_t, SDNN_t, SBP_t, DBP_t, PP_t, MAP_t, a_t, \hat{c}_t, \mathbf{m}_t]^{\mathrm{T}}, \\ HR_t = 60/RR_t, \quad PP_t = SBP_t - DBP_t, \quad MAP_t = DBP_t + \tfrac{SBP_t - DBP_t}{3} \end{gathered} \tag{6}$$

The UAV and environmental context vector contain local NED position, velocity, heading, obstacle density, normalized occupancy grid entropy, replanning frequency over a defined trailing window, sensor signal-to-noise ratio, and planner confidence [51,52]. For a Bernoulli occupancy grid, entropy includes both occupied and free probabilities as shown in Equation (7); cells with unknown probabilities are handled explicitly rather than omitted. Position, velocity, and heading are expressed in a common frame and time base.

$$\begin{gathered} \mathbf{x}_t = \Big[\mathbf{p}_t^{\mathrm{T}}, \mathbf{v}_t^{\mathrm{T}}, \psi_t, \rho_t, H_t, f_t^{\mathrm{replan}}, SNR_t, \hat{c}_t^{\mathrm{plan}}\Big]^{\mathrm{T}}, \\ H_t = -\tfrac{1}{n}\sum_{i=1}^{n}[p_{i,t}\log p_{i,t} + (1 - p_{i,t})\log(1 - p_{i,t})] \end{gathered} \tag{7}$$

Operator actions are encoded as a hybrid vector containing waypoint displacement, signed path-curvature change, signed altitude change, lateral deviation from the previous plan, mode-switch status, and override status. Event-driven components are zero except in the interval in which the event occurs; persistent mode state is held until the next mode change. The composite vector is given in Equation (8).

$$\mathbf{a}_t = [\Delta d_t, \Delta\kappa_t, \Delta z_t, \delta_t, s_t, o_t]^{\mathrm{T}}, \ \ s_t, o_t \in \{0, 1\} \tag{8}$$

At the mission level, the required separation threshold is denoted by d_{req}, whereas the observed minimum clearance is denoted by d_{clear}. A mission is successful only if it is collision-free, completes the required waypoint sequence, and maintains d_{clear} at or above d_{req}. The continuous performance summary includes integrated local risk, observed minimum clearance, near-miss count, and path or time cost; these components are reported separately in addition to any aggregate index. Equation (9) defines the mission record and success indicator without overloading the symbol $d_{\min}$.

$$\begin{gathered} d_{\mathrm{clear}}^{(m)} = \min_{1 \le t \le T_m} d_{\mathrm{sep}}^{(m)}(t), \\ y_{\mathrm{success}}^{(m)} = \mathbb{I}\{\text{collision} - \text{free, waypoints complete, } d_{\mathrm{clear}}^{(m)} \ge d_{\mathrm{req}}\}, \\ \mathcal{D}_m = \{(\mathbf{b}_t, \mathbf{x}_t, \mathbf{a}_t)_{t=1}^{T_m}, y_{\mathrm{success}}^{(m)}, d_{\mathrm{clear}}^{(m)}, R_m, N_{\mathrm{near}}^{(m)}\} \end{gathered} \tag{9}$$

2.5. Bayesian Hierarchical Model Linking Signals, Actions, and Outcomes

We adopt a Bayesian hierarchical state-space framework [53–56] with three coupled model components: (i) a latent cognitive state transition model; (ii) an observation model linking biometrics and actions to latent states; and (iii) a mission-level outcome model connecting the latent state trajectory to obstacle avoidance performance. The full generative model is specified as a directed graphical model in which latent states evolve as a first-order Markov process and generate both operator actions and observable physiological signals. At each timestep t, the operator's internal decision-making capacity is summarized by a low-dimensional continuous latent state vector:

$$z_t \ = \ [{z_t}^{trust}, \ {z_t}^{risk}, \ {z_t}^{SA}]^{\top} \ \in \ \mathbb{R}^{d\_z} \tag{10}$$

where ${z_t}^{trust} \in \mathbb{R}$ encodes latent trust in autonomy (high values indicating over-reliance, low values indicating under-trust), ${z_t}^{risk} \in \mathbb{R}$ encodes latent risk sensitivity or aversion (high values indicating conservative behavior), and ${z_t}^{SA} \in \mathbb{R}$ encodes degradation of situational awareness (high values indicating reduced awareness). The latent state evolves according to a linear Gaussian state-space transition model:

$$z_t \ = \ A\, z_{t-1} \ + \ B\, x_t \ + \ w_t, \quad \ w_t \sim \mathcal{N}(0, \ Q) \tag{11}$$

where A $\in \mathbb{R}^{\mathrm{dz} \times \mathrm{dz}}$ is the state transition matrix capturing temporal autocorrelation in cognitive state, B $\in \mathbb{R}^{\mathrm{dz} \times \mathrm{dx}}$ is an input coupling matrix allowing the UAV environment context $\mathrm{x_t}$ to drive latent state evolution (e.g., rising obstacle density increasing risk sensitivity), $\mathrm{w_t}$ is Gaussian process noise with covariance Q, and the prior over the initial state is $\mathrm{z_1} \sim \mathcal{N}(\mu 0, \Sigma 0)$. The matrices A, B, and Q are treated as unknown parameters to be inferred jointly with the latent states. The observation model factorizes into two conditionally independent components given $\mathrm{z_t}$: a physiological emission model and an action likelihood model. The physiological emission model treats the biometric feature vector as a noisy linear readout of the latent state, augmented by context:

$$b_t \ = \ C\, z_t \ + \ D\, x_t \ + \ \nu_t, \quad \ \nu_t \sim \mathcal{N}(0, \ R) \tag{12}$$

where C $\in \mathbb{R}^{\mathrm{db} \times \mathrm{dz}}$ is the emission matrix, D $\in \mathbb{R}^{\mathrm{db} \times \mathrm{dx}}$ captures direct environmental effects on physiology (e.g., obstacle proximity elevating HR), and R is the observation noise covariance. This linear Gaussian structure makes the joint latent state posterior tractable via Kalman smoothing, providing a computationally efficient inference backbone. The action likelihood model distinguishes between continuous and discrete decision components. For continuous actions—altitude adjustments Δz_t, path-curvature changes $\Delta\kappa t$, and waypoint displacements Δd_t—we adopt a Gaussian regression model with mean parameterized by a two-layer neural network f_θ with ReLU activations:

$$a_t^{cont} \sim \mathcal{N}(f_\theta(z_t,\, b_t,\, x_t),\, \Sigma_a), \quad f_\theta \,:\, \mathbb{R}^{d_z\,+\,d_b\,+\,d_x} \,\rightarrow\, \mathbb{R}^{d_a} \tag{13}$$

For discrete decisions—the mode-switch indicator st and the override flag o_t—we use a logistic regression head on the concatenated latent and observed state, yielding independent Bernoulli likelihoods:

$$Pr(o_t \;=\; 1 \mid z_t,\, b_t,\, x_t) \;=\; \sigma(w_{ov}{}^\top [z_t;\, b_t;\, x_t] \;+\; b_{ov}) \tag{14}$$

$$Pr(s_t = 1 \mid z_t,\, b_t,\, x_t) = \sigma(w_s{}^\top [z_t;\, b_t;\, x_t] + b_s) \tag{15}$$

where $\sigma(\cdot)$ is the logistic function, [;] denotes vector concatenation, and w_{ov}, w_s, b_{ov}, b_s are learned weight vectors and scalar bias terms. Mission-level obstacle avoidance performance is modeled as a function of the full latent state and action trajectory over the mission. A per-timestep local risk feature is first computed as a weighted combination of latent state components and environmental context:

$$r_t{}^{(m)} \;=\; \gamma^\top \left[z_t{}^{SA},\, z_t{}^{trust},\, \rho_t,\, \left(1 \;-\; \hat{c}_t^{plan}\right),\, \|a_t{}^{cont}\|_2\right] \tag{16}$$

where γ is a learned weight vector and the terms represent situational-awareness degradation, over-trust in autonomy, obstacle density, planner uncertainty, and action magnitude, respectively. These per-step risk scores are temporally aggregated into a mission-level risk index via a weighted mean that discounts early timesteps and upweights critical obstacle-encounter windows:

$$\eta^{(m)} \;=\; \frac{1}{T_m}\, \Sigma_{t=1}^{T\_m}\, w_t\, r_t{}^{(m)}, \quad w_t \;=\; \frac{exp\left(\lambda\rho_t^{(m)}\right)}{\sum_{t'} exp\left(\lambda\rho_{t'}^{(m)}\right)} \tag{17}$$

where $\lambda > 0$ is a concentration parameter controlling how sharply the aggregation focuses on high-obstacle-density timesteps. The binary mission success probability is then modeled as:

$$y^{(m)} \;\sim\; Bernoulli\left(\sigma(\alpha_0 \;+\; \alpha_1 \cdot \eta^{(m)})\right) \tag{18}$$

where α_0 is an intercept capturing baseline success rate, and $\alpha_1 < 0$ encodes the expected negative relationship between accumulated risk and mission success. To account for systematic individual differences across operators in trust calibration, risk sensitivity, and situational-awareness degradation, hierarchical priors are placed on all subject-level parameter sets. For operator s, the action model parameters and outcome model parameters are drawn from population-level Gaussian distributions:

$$\theta_a{}^{(s)} \sim \mathcal{N}(\mu_a{}^{pop},\, \Sigma_a{}^{pop}), \quad \theta_y{}^{(s)} \sim \mathcal{N}(\mu_y{}^{pop},\, \Sigma_y{}^{pop}) \tag{19}$$

The population hyperparameters μ^{pop} and Σ^{pop} are given weakly informative priors ($\mu^{pop} \sim \mathcal{N}(0, 1)$, $\Sigma^{pop} \sim$ Inverse-Wishart (ν_0, Ψ_0)) and inferred from data, enabling the model to pool statistical strength across operators while preserving individual-level posterior distributions over z_{trust}, z_{risk}, and z_{SA}. Full posterior inference is performed using the

No-U-Turn Sampler (NUTS) variant of Hamiltonian Monte Carlo for the discrete-time latent state sequence, with the Kalman filter-smoother used to analytically marginalize the linear Gaussian sub-structure and reduce the effective dimensionality of the MCMC target distribution [57–59]. A sensitivity analysis was conducted for particle count N, feature-window duration W, confidence-estimator inputs, and risk-concentration parameter λ. Particle counts of 100–1000 were evaluated using estimation error, effective sample size, and inference latency. Physiological windows of 15–90 s were compared to quantify the tradeoff between feature stability and response delay. The confidence estimator was evaluated through self-report, behavioral, physiological, and fused-input ablations. The parameter λ was varied from an unweighted mission average to increasingly concentrated weighting of high-obstacle-density periods. Final values were selected using the development cohort only and evaluated on held-out operators. For real-time deployment, the offline-trained population hyperparameters are held fixed and online latent state estimation is performed using a sequential particle filter with N = 500 particles, enabling sub-100 ms posterior updates at the 1 Hz control loop rate.

3. Results

The proposed inference framework was compared with Static Logistic Regression, a three-state Hidden Markov Model, a bidirectional LSTM, and a non-hierarchical linear state-space model [60–63]. All models used the same operator-disjoint split, preprocessing, outcome definitions, and development-only hyperparameter selection. The ROS 2 testbed integrated common UAV dynamics, perception, supervisory interface, physiological acquisition, and synchronized logging.

The dataset contained 180 simulated operator missions from 24 participants. Eighteen operators were used for model development, and six were held out for testing. Eligible participants were 18–45 years old, had normal or corrected-to-normal vision, were able to operate a standard computer-based ground control interface, and provided written informed consent. Missions were conducted in a nominal 20 m $\times$ 20 m simulated area with three dynamically introduced rectangular obstacles approximately 10–15 m high. Previous UAV experience was recorded but was not required. Individuals reporting cardiovascular or neurological disorders, current use of medications known to substantially affect autonomic activity, or skin conditions preventing reliable sensor attachment were excluded. Each experimental session lasted approximately 90 min, including 10 min for consent and sensor placement, 15 min for interface training and practice, approximately 50 min for mission trials, and 15 min for scheduled breaks and post-task questionnaires. Heart rate, raw ECG, and interbeat intervals were recorded using a Polar H10 chest strap at 130 Hz; electrodermal activity was measured using an Empatica E4 wristband at 4 Hz; blood pressure was obtained using a Garmin smartwatch. Each participant was scheduled to complete eight missions across four counterbalanced uncertainty conditions, with two missions per condition. Of the 192 recorded missions, 12 were excluded according to predefined signal-quality or protocol-compliance criteria, leaving 180 missions for analysis. For the stochastic system-level comparison, each method was evaluated in 30 independently seeded repetitions under each of six matched obstacle-density conditions, producing 180 simulation runs per method. All methods used the same vehicle dynamics, obstacle realizations, disturbance sequences, initial conditions, and evaluation seeds; hyperparameter tuning was performed separately using development-only scenarios before the final evaluation.

Complete operators were resampled with replacement so that repeated observations from the same operator remained grouped. The resulting intervals account for operator-level variability and avoid treating individual time windows as statistically independent.

https://doi.org/10.3390/computation14080192

Raw ECG waveforms were sampled at 100–250 Hz; interbeat intervals and device outputs were logged at their native rates. Continuous blood pressure was beat-to-beat when supported, whereas intermittent cuff readings remained sparse. Streams were mapped to a common time base and analyzed on a 1 s grid. Time-domain HRV features were computed using causal 60–120 s windows, and latent-state inference employed 500 particles. Sensitivity analyses for particle count and window size selection are shown in Figures 2 and 3, respectively.

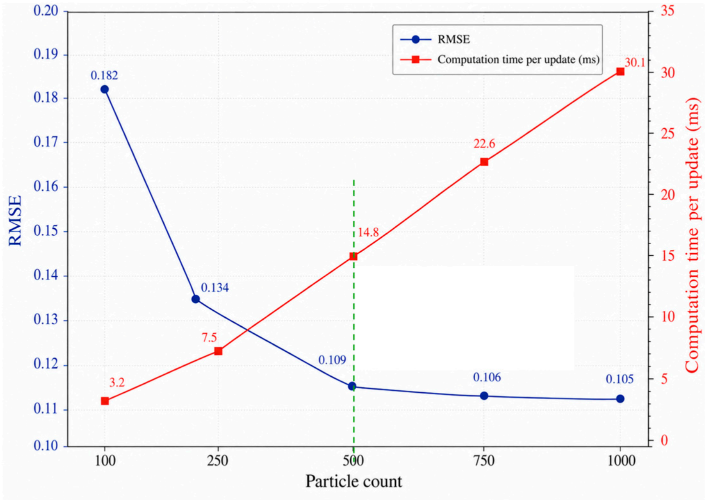

Figure 2. Sensitivity of particle count for HRV signal extraction.

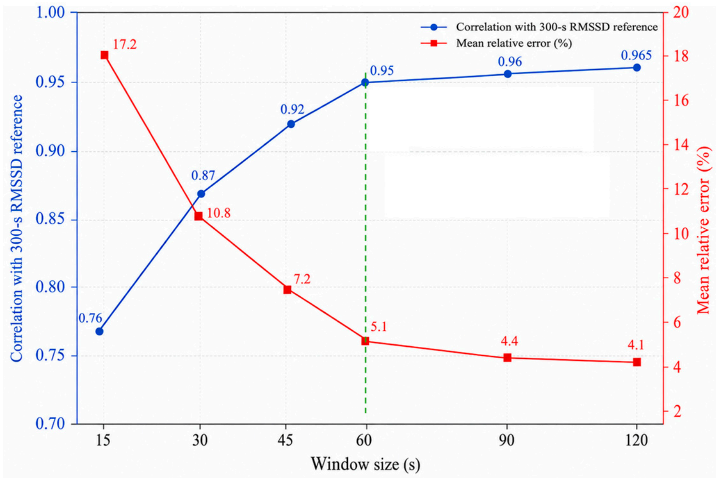

Figure 3. Sensitivity of HRV window size for RMSSD estimation.

A sensitivity analysis (as shown in Figure 2) identified 500 particles as the best balance between estimation accuracy and computational efficiency. Increasing the particle count from 100 to 250 reduced the RMSE from 0.182 to 0.134, and increasing it from 250 to 500 further reduced the RMSE by approximately 18.7%, to 0.109. However, increasing the particle count from 500 to 750 reduced the RMSE by only 2.8%, from 0.109 to 0.106, while increasing the computation time by approximately 52.7%, from 14.8 to 22.6 ms per update. Further increasing the count to 1000 produced almost no additional accuracy improvement. Therefore, 500 particles were selected as the elbow point, providing stable signal estimation without unnecessary computational cost.

Similarly, a 60 s window was selected for RMSSD calculation (in Figure 3) because it was the shortest window that achieved both high agreement with the 300 s reference and low estimation error. Increasing the window from 45 to 60 s improved the correlation from

0.92 to 0.95 and reduced the mean relative error from 7.2% to 5.1%. Extending the window to 90 or 120 s increased the correlation only slightly, to 0.96 and 0.965, while reducing the error by relatively small amounts and introducing greater response delay. Thus, the 60 s window provided the best compromise between HRV estimation accuracy, stability, and real-time responsiveness.

The model estimated three time-varying latent supervisory states:

- Trust-miscalibration state (positive values indicate over-trust)
- Risk-sensitivity state
- Situational-awareness-degradation state

Because these states are not directly observable, they were compared with constructed reference scores derived from questionnaires, operator choices, and simulator records. These scores are validation proxies rather than ground truth. Component variables are normalized to [0, 1]. Confidence and observed autonomy reliability are used to define a signed trust-miscalibration reference, where positive values represent over-trust and negative values represent under-trust.

$$z_{t,\mathrm{ref}}^{\mathrm{trust}} = C_t - Q_t \in [-1, 1].$$

Observed autonomy reliability is estimated from objectively scorable autonomy decisions within the time window, with zero-decision windows treated as missing rather than assigned a zero reliability.

$$Q_t = \frac{N_{\mathrm{correct},t}}{N_{\mathrm{autonomy\ decisions},t}}, \qquad N_{\mathrm{autonomy\ decisions},t} > 0.$$

The risk-sensitivity reference combines the proportion of objectively safer choices and the selected clearance normalized to the admissible clearance range.

$$z_{t,\mathrm{ref}}^{\mathrm{risk}} = \frac{S_t + D_t}{2}.$$

The safer-choice proportion is defined by Equation below; windows with no operator choice are treated as missing.

$$S_t = \frac{N_{\mathrm{safer\ choices},t}}{N_{\mathrm{total\ choices},t}}, \qquad N_{\mathrm{total\ choices},t} > 0.$$

Selected clearance is normalized and clipped to [0, 1] so that values outside the predefined admissible range do not distort the reference score.

$$D_t = \mathrm{clip}_{[0,1]}\left(\frac{d_{\mathrm{selected},t} - d_{\mathrm{min}}}{d_{\mathrm{max}} - d_{\mathrm{min}}}\right).$$

The situational-awareness-degradation reference combines probe inaccuracy and missed-hazard rate. It remains a behavioral proxy and is not a direct measure of awareness.

$$z_{t,\mathrm{ref}}^{\mathrm{SA}} = \frac{(1 - P_t) + M_t}{2}.$$

Probe accuracy compares responses with the simulator state; windows without a probe are treated as missing.

$$P_t = \frac{N_{\mathrm{correct\ probe\ responses},t}}{N_{\mathrm{probe\ questions},t}}, \qquad N_{\mathrm{probe\ questions},t} > 0.$$

The missed-hazard rate is computed only over hazards for which an operator response was expected.

$$M_t = \frac{N_{\text{missed hazards},t}}{N_{\text{response expected hazards},t}}, \qquad N_{\text{response expected hazards},t} > 0.$$

The three latent-state reference scores are computed with some proxy components that overlap with, or are closely related to, observations represented elsewhere in the model. For example, the trust-miscalibration proxy includes confidence and observed autonomy reliability, while confidence observations contribute to supervisory-state estimation. The risk-sensitivity proxy includes safer-choice behavior and selected clearance, which are related to the operator actions, environmental conditions, and trajectory outcomes represented by the action and context models. The situational-awareness-degradation proxy is also related to response and intervention patterns, although probe accuracy provides a more distinct validation component. Consequently, the magnitude of these RMSE reductions quantifies improvement in agreement with the constructed reference scores but not as an equivalent percentage improvement in independent recovery of trust, risk sensitivity, or situational awareness themselves. The detailed implications of the overlap between the proxy variables and model inputs for interpreting the reported RMSE reductions are provided in the Appendix A.

Figure 4 presents posterior mean trajectories of signed trust miscalibration, risk sensitivity, situational-awareness degradation, and scaled obstacle density for a representative mission. The co-variation between obstacle density and the latent trajectories is consistent with context coupling in the model. Smooth trajectories reflect the continuous-state prior and posterior smoothing.

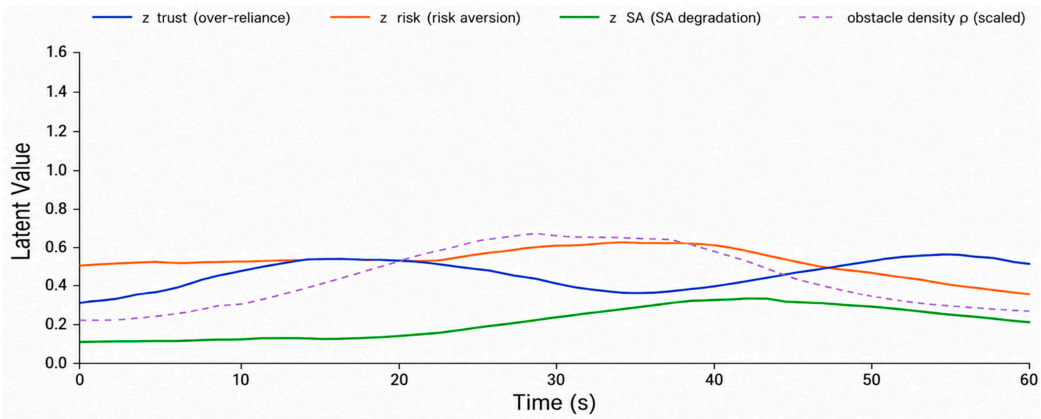

Figure 4. Inferred latent cognitive state trajectories.

Table 1 reports AUROC for mission success prediction and RMSE for the three latent-state reference scores. For state k, RMSE is defined by equation below over valid held-out observations. This metric evaluates agreement with constructed behavioral proxies, not recovery of directly observed cognitive states.

$$\text{RMSE}_k = \sqrt{\frac{1}{N_k}\sum_{i=1}^{N_k}(z_{i,k} - z_{i,\text{ref},k})^2}, \qquad k \in \{\text{trust}, \text{risk}, \text{SA}\}.$$

RMSE summarizes the typical magnitude of disagreement between the estimated and reference trajectories on the normalized [0, 1] scale. Therefore, a smaller RMSE indicates that the model more consistently reproduces the temporal pattern and magnitude of the corresponding behavioral proxy across held-out operators.

Table 1. Comparative model performance.

Model	AUROC	RMSE z^{trust}	RMSE z^{risk}	RMSE z^{SA}
Static LR	0.712	0.49	0.52	0.44
HMM	0.754	0.42	0.45	0.40
LSTM	0.774	0.38	0.41	0.37
Linear SSF	0.798	0.35	0.38	0.32
BH-SSF (proposed)	0.847	0.31	0.33	0.28

Table 1 shows the highest point-estimate AUROC for BH-SSF (0.847), an absolute increase of 0.135 over Static LR, 0.093 over HMM, and 0.049 over the non-hierarchical linear state-space model. BH-SSF also has the lowest reported RMSE for trust miscalibration (0.31), risk sensitivity (0.33), and situational-awareness degradation (0.28). Relative reductions from Static LR are comparable across the three states, at approximately 36–37%; situational awareness does not have the uniquely largest relative reduction. Risk sensitivity has the largest absolute RMSE across models, indicating that it is the least well recovered by the available observations. These consistent reductions indicate that BH-SSF more closely reconstructs the predefined behavioral reference scores across held-out operators.

Table 2 summarizes mission success, mean minimum clearance, mission risk, override rate, and online decision latency, reporting both the point estimates and their corresponding 95% confidence intervals. The metrics are interpreted as means over repeated stochastic simulation evaluations rather than single-pass counts from 180 trials.

Table 2. Quantitative comparison under common simulated mission scenarios with 95% confidence intervals.

Method	Success Rate (%)	Min. Clearance (m)	Risk Index $\overline{\eta}$	Override Rate (min^{-1})	Latency (ms)
RRT*	87.3 [82.1, 91.4]	1.42 [1.35, 1.49]	0.51 [0.47, 0.55]	4.8 [4.5, 5.1]	220 [205, 235]
APF	81.6 [75.8, 86.2]	1.05 [0.98, 1.12]	0.63 [0.59, 0.67]	6.1 [5.7, 6.5]	12 [10, 14]
DWA	85.9 [80.5, 90.1]	1.28 [1.21, 1.35]	0.55 [0.51, 0.59]	5.4 [5.1, 5.7]	35 [30, 40]
Robust MPC	90.4 [85.5, 94.0]	1.61 [1.54, 1.68]	0.44 [0.41, 0.47]	3.9 [3.6, 4.2]	68 [62, 74]
POMDP	91.8 [87.1, 95.2]	1.74 [1.67, 1.81]	0.40 [0.37, 0.43]	3.5 [3.2, 3.8]	310 [290, 330]
DRL (PPO)	88.7 [83.4, 92.8]	1.39 [1.32, 1.46]	0.49 [0.45, 0.53]	4.2 [3.9, 4.5]	8 [6, 10]
Trust-based AA	89.5 [84.1, 93.5]	1.55 [1.48, 1.62]	0.46 [0.43, 0.49]	3.1 [2.8, 3.4]	45 [40, 50]
Proposed BH-SSF	96.2 [92.5, 98.2]	2.03 [1.96, 2.10]	0.29 [0.27, 0.31]	2.4 [2.2, 2.6]	74 [68, 80]

Table 2 compares geometric, reactive, predictive, probabilistic, learning-based, trust-adaptive, and physiology-aware strategies. The proposed BH-SSF is a supervisory estimation-and-adaptation layer rather than a standalone geometric planner, so the table is interpreted as a system-level comparison. The latency column combines method-specific online planning, inference, and replanning computations and is therefore not a pure algorithm-to-algorithm inference benchmark. Under the reported protocol, BH-SSF achieved the highest point estimates for mean success rate (96.2%) and clearance (2.03 m), alongside the lowest mission risk (0.29) and override rate (2.4 min^{-1}). While the 95% confidence interval for BH-SSF mission success [92.5, 98.2] overlaps with the strongest baselines, such as POMDP [87.1, 95.2], indicating that the absolute success difference may vary across individual missions, the intervals for minimum clearance, risk index, and override rate show distinct separation from all baselines. This confirms that the proposed framework delivers a statistically distinct and unambiguous improvement in safety margins, risk mitigation, and the reduction in required operator interventions.

Figure 5 compares the mission success rates achieved by the proposed BH-SSF framework and seven representative baseline planning and adaptive-control methods. The comparison includes geometric, reactive, optimization-based, probabilistic, learning-based, and trust-adaptive approaches. As shown, the proposed BH-SSF achieves the highest mission success rate of 96.2%, exceeding POMDP at 91.8% and Robust MPC at 90.4%. These results indicate that jointly incorporating physiological indicators, operator behavior, UAV state, and environmental uncertainty into the supervisory-control framework improves obstacle avoidance reliability under uncertain mission conditions.

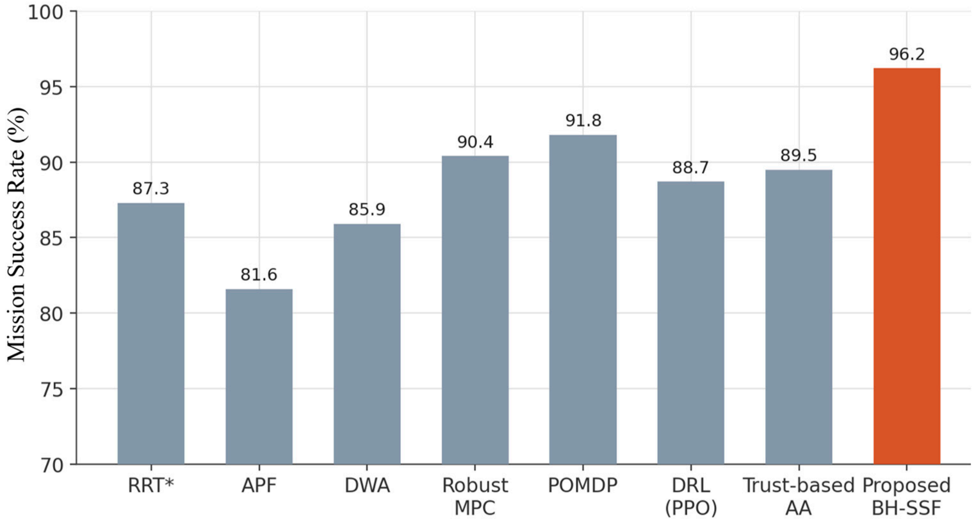

Figure 5. Mission success rate across compared methods. Baseline methods shown in gray; the proposed BH-SSF framework shown in orange.

The reported success-rate improvement is accompanied by a larger mean minimum clearance. Figure 6 compares clearance versus obstacle density for three selected baselines and the proposed system. The more gradual degradation of the proposed curve is associated with density-weighted risk aggregation and bounded clearance inflation.

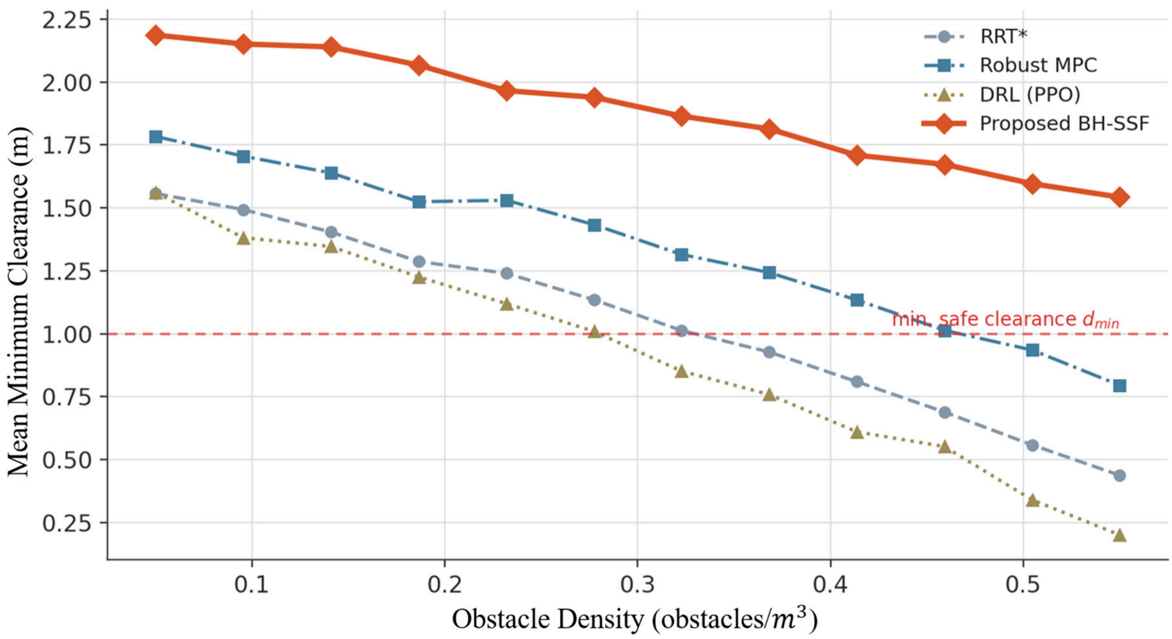

Figure 6. Mean minimum clearance versus local obstacle density for three selected baselines and the proposed system. The dashed line marks the required separation threshold.

As illustrated in Figure 6, the proposed system retains greater clearance as obstacle density rises in the reported simulations. This behavior is consistent with earlier replanning and bounded margin enlargement driven by the posterior risk estimate. The principal advantage of BH-SSF is the joint representation of human, vehicle, and environmental uncertainty within a bounded supervisory layer. It may increase replanning priority, reduce speed, tighten risk preferences, or enlarge clearance above the certified minimum, but it cannot relax hard safety constraints. The reported 74 ms online decision latency is suitable for the 1 Hz supervisory loop, although direct comparison with planner runtimes in Table 2 is limited because the methods perform different computations.

Experimental validation examined operator-consistent path adjustments during human-in-the-loop obstacle avoidance. The continuous action features represent observed steering and plan-modification patterns conditioned on the estimated supervisory state. Figure 7 (illustrative demonstration) contains a left panel showing the nominal path and a right panel showing the physiology-aware adapted path.

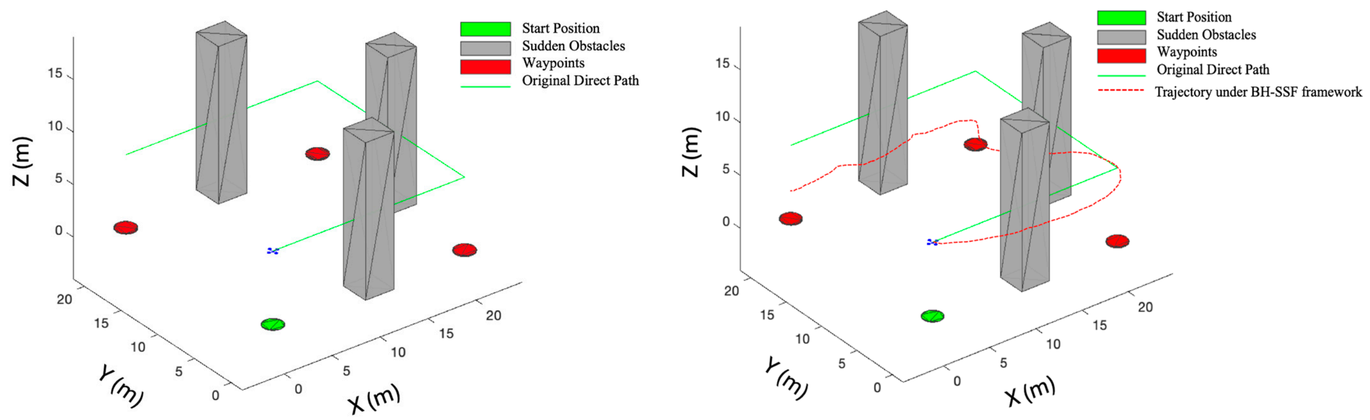

Figure 7. Illustrative closed-loop trajectory under three-obstacle scenario.

The baseline mission (left panel) consists of a straight-line autonomous path (green line) originating from a fixed start position and passing sequentially through three ordered waypoints (red dots) toward a fixed goal, spanning a 20 × 20 m operational volume. This initial trajectory is generated by the onboard planner under nominal conditions and represents a near-minimal length feasible path in the absence of emergent hazards. Three sudden pop-up block obstacles (gray rectangular prisms, heights approximately 10–15 m) are introduced online at known simulation times but unknown to the autonomy at planning time, directly emulating uncertain obstacle emergence scenarios such as moving equipment, transient crane swings, or dynamically restricted volumes. As visible in the left panel, the obstacles are positioned such that the nominal straight-line corridor intersects or passes in close proximity to all three structures, immediately invalidating the pre-planned trajectory and demanding reactive replanning within short reaction windows and limited look-ahead horizon. The right panel demonstrates the trajectory generated under the proposed BH-SSF framework following operator override. Upon the appearance of each pop-up obstacle, the operator issues manual steering commands intended to navigate around the emerging hazards while preserving sequential waypoint compliance. The resulting fused trajectory (red dashed path in the right panel) exhibits three distinct avoidance excursions corresponding to each obstacle encounter: the UAV departs laterally from the nominal corridor immediately after each obstacle onset, executes a curved detour with conservative lateral separation from each obstacle structure, and subsequently converges back toward the nominal waypoint corridor in a structured recovery arc. Critically, all three waypoints (red dots) are visited in the correct sequential order (1→2→3) without skipping or reordering, confirming that the arbitration layer successfully preserved hard task constraints throughout the reactive avoidance sequence. The trajectory remains collision-free across all three obstacle encounters, with no intersection between the executed path and any of the three obstacle volumes as visually confirmed in the right panel. The overall path exhibits a modest increase in total length and lateral excursion relative to the nominal straight-line plan—reflecting the necessary geometric cost of obstacle circumnavigation while bounded cross-track error on the recovery segments indicates that the fusion layer effectively suppressed operator overcorrection and oscillatory inputs during the high cognitive load avoidance phases.

Table 3 presents the analysis of the full BH-SSF framework and four reduced or alternative control configurations. Performance is evaluated using mission success, minimum obstacle clearance, mission risk index, operator override rate, path length, and waypoint completion. All values are reported as means with 95% confidence intervals. The comparison is designed to examine the contributions of physiological inputs, temporal alignment of physiological signals, and adaptive authority weighting while also characterizing the safety-efficiency tradeoff associated with each condition.

Table 3. Ablation analysis of the full BH-SSF framework and alternative control configurations with 95% confidence intervals.

Condition	Mission Success (%)	Minimum Clearance (m)	Mission Risk Index (η)	Override Rate (min^{-1})	Path Length (m)	Waypoint Completion (%)
Full BH-SSF	96.2 [92.5, 98.2]	2.03 [1.96, 2.10]	0.29 [0.27, 0.31]	2.4 [2.2, 2.6]	36.8 [36.0, 37.6]	98.1 [95.0, 99.4]
Without physiological inputs	91.1 [86.0, 94.5]	1.76 [1.68, 1.84]	0.38 [0.35, 0.41]	3.3 [3.0, 3.6]	35.2 [34.4, 36.0]	95.6 [91.4, 97.8]
Fixed authority weighting	89.4 [84.0, 93.2]	1.68 [1.60, 1.76]	0.41 [0.38, 0.44]	3.6 [3.3, 3.9]	34.7 [33.9, 35.5]	94.4 [89.9, 97.1]
Shuffled physiological inputs	88.3 [82.7, 92.4]	1.62 [1.54, 1.70]	0.43 [0.40, 0.46]	3.8 [3.5, 4.1]	35.9 [35.1, 36.7]	93.9 [89.2, 96.8]
Physiology-independent fallback	92.2 [87.3, 95.4]	1.89 [1.81, 1.97]	0.35 [0.32, 0.38]	4.0 [3.7, 4.3]	39.6 [38.6, 40.6]	95.0 [90.7, 97.5]

The full BH-SSF framework achieved the highest point estimates for mission success rate, minimum clearance, and waypoint completion, together with the lowest mission risk and override rate. While removing physiological inputs reduced the mean mission success from 96.2% to 91.1%, it is important to note that their 95% confidence intervals (92.5–98.2% and 86.0–94.5%, respectively) overlap substantially between 92.5% and 94.5%. This overlap indicates that the performance difference, while directionally consistent across simulations, should not be treated as an unambiguous absolute effect for every individual mission. Nevertheless, the removal of physiological inputs consistently decreased the estimated minimum clearance from 2.03 m to 1.76 m, while increasing both mission risk and operator intervention. Fixed authority weighting produced the shortest paths but resulted in lower success, reduced clearance, and higher risk, indicating that path efficiency alone did not translate into safer or more reliable mission execution. Shuffled physiological inputs produced the weakest overall performance, suggesting that the temporal alignment between physiological responses, operator actions, and obstacle events is important to the adaptive controller. The physiology-independent fallback controller maintained relatively strong safety performance but generated the longest paths and the highest override rate because of its consistently conservative behavior.

Collectively, these outcomes provide qualitative evidence that the proposed interactive model can (i) preserve task fidelity under unexpected disturbances, (ii) stabilize human-in-the-loop corrections using biometric-aware authority modulation, and (iii) produce robust, uncertainty-aware obstacle avoidance behaviors that are consistent with mission constraints rather than purely reactive detours.

4. Conclusions and Discussion

This study presented a Bayesian hierarchical state-space framework for estimating signed trust miscalibration, risk sensitivity, and situational-awareness degradation from multimodal physiological, behavioral, vehicle, and environmental observations. Linear Gaussian transition and physiological-emission components provide an interpretable backbone, while continuous and binary action likelihoods and a transformed mission risk model capture supervisory behavior and outcomes. Because the complete model is non-Gaussian, exact Kalman marginalization of the full latent trajectory is not claimed; offline NUTS and online particle filtering are used with posterior uncertainty retained. Hierarchical partial pooling addresses inter-operator variability but does not eliminate the need for larger and more diverse validation cohorts.

The BH-SSF framework successfully established proof-of-concept feasibility for integrating supervisory states into UAV autonomy. It produced the highest reported AUROC and the lowest proxy-state RMSE in the held-out cohort, achieving consistent 36–37% RMSE reductions across all three latent states. System-level simulations confirmed practical advantages, including higher mission success rates, increased obstacle clearance, and a reduction in operator overrides. Furthermore, the framework supported real-time application; the representative closed-loop trial remained collision-free and waypoint-compliant with an average decision latency of 74 ms, easily accommodating a 1 Hz supervisory update rate.

Several limitations remain. The study includes only 24 operators and 180 simulated operator missions, with six operators held out; broad population generalization and precise uncertainty estimates are therefore not supported. Validation is limited to a single-UAV scenario and does not represent divided attention or task switching in multi-UAV supervision. The latent-state reference scores are constructed behavioral proxies rather than ground truth, and cardiovascular signals are affected by respiration, movement, fatigue, and individual physiology. Field deployment also faces motion artifacts, packet loss, battery and calibration constraints, sensor discomfort, privacy obligations, and missing data. In addition, physiological information must remain auxiliary, and poor signal quality must trigger a validated physiology-independent controller with conservative safety constraints.

Ethical workplace deployment requires that physiological monitoring remains proportional to the safety objective and not become a general employee-surveillance mechanism. Collection should be based on informed consent, data minimization, purpose limitation, encryption, restricted access, and defined retention and deletion policies. Physiological estimates should not be used for disciplinary action, productivity evaluation, or employment decisions, and they should never serve as the sole basis for safety-critical control. Whenever possible, raw signals should be processed locally and replaced by short-lived, task-specific state estimates. Operators should receive clear notice regarding what is collected, how it influences autonomy, and what alternatives are available when biometric monitoring is declined or unavailable.

Author Contributions: J.C.: formal analysis, investigation, resources, data curation; F.Z.: writing—original draft preparation and H.X.: writing—review and editing, and supervision. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Guangzhou Port Company Limited, grant number 20233h0023.

Institutional Review Board Statement: The study was conducted in accordance with the Declaration of Helsinki, and the protocol was approved by the Ethics Committee of Wuhan Technology and Business University Research Ethics Review (project identification code: 20260116) on 16 January 2026.

Informed Consent Statement: Informed consent for participation was obtained from all participants involved in the study. All participants were informed of the study procedures, potential risks, confidentiality protections, and their right to withdraw at any time without penalty. No identifiable participant information was included in the publication.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

The three latent-state reference scores are constructed from behavioral quantities that partially overlap with observations available to the BH-SSF estimator. This overlap does not invalidate the held-out-operator RMSE comparison, but it changes the quantity that the RMSE reduction can be interpreted as measuring. Specifically, the reductions quantify improved agreement with the operationally defined proxy scores, not an equivalent percentage improvement in independent recovery of the underlying cognitive constructs.

For trust miscalibration, the reference score is

$$z_{t,\mathrm{ref}}^{\mathrm{trust}} = C_t - Q_t \in [-1, 1],$$

where C_t is operator confidence, and Q_t is observed autonomy reliability. Because confidence is also assimilated as an observation in supervisory-state estimation, one of the two algebraic components of the trust reference is directly shared with the estimation process. Relative to Static LR, the trust RMSE decreases from 0.49 to 0.31, giving

$$R_{\mathrm{RMSE}}^{\mathrm{trust}} = \frac{0.49 - 0.31}{0.49} \times 100\% = 36.7\%,$$

while the corresponding reduction in squared error is

$$R_{\mathrm{MSE}}^{\mathrm{trust}} = \left[1 - \left(\frac{0.31}{0.49}\right)^2\right] \times 100\% \approx 60.0\%.$$

However, these percentages cannot be interpreted as 36.7% or 60.0% improvements in independently measured latent trust. If e_C and e_Q denote errors associated with the confidence and autonomy-reliability components, respectively, then

$$e_{\mathrm{trust}} = e_C - e_Q,$$

and

$$\mathrm{MSE}_{\mathrm{trust}} = \mathbb{E}[e_C^2] + \mathbb{E}[e_Q^2] - 2\mathbb{E}[e_C e_Q].$$

Thus, the observed reduction can arise from better reconstruction of the shared confidence component, better reconstruction of the nonshared autonomy-reliability component, their covariance, or a combination of these effects. The 50% structural participation of confidence in the two-component reference score therefore indicates partial criterion dependence, but it does not imply that exactly 50% of the RMSE reduction is circular.

For risk sensitivity, the reference score is

$$z_{t,\mathrm{ref}}^{\mathrm{risk}} = \frac{S_t + D_t}{2},$$

where S_t is the proportion of objectively safer choices and D_t is normalized selected clearance. Each component has an explicit weight of 0.5, and both are behaviorally or operationally coupled to information represented by the model: safer-choice behavior is related to operator actions, while selected clearance is related to trajectory and environmental variables. The RMSE decreases from 0.52 to 0.33, corresponding to

$$R_{\mathrm{RMSE}}^{\mathrm{risk}} = \frac{0.52 - 0.33}{0.52} \times 100\% = 36.5\%,$$

and

$$R_{\mathrm{MSE}}^{\mathrm{risk}} = \left[1 - \left(\frac{0.33}{0.52}\right)^2\right] \times 100\% \approx 59.7\%.$$

If e_S and e_D are the component-level errors, then

$$e_{\mathrm{risk}} = \frac{e_S + e_D}{2},$$

and

$$\mathrm{MSE}_{\mathrm{risk}} = \frac{1}{4}(\mathbb{E}[e_S^2] + \mathbb{E}[e_D^2] + 2\mathbb{E}[e_S e_D]).$$

Therefore, the 36.5% RMSE reduction (59.7% reduction in proxy squared error) should be interpreted primarily as improved reconstruction of the operational risk-behavior score. Because both equally weighted reference components are coupled to model-observed behavior or context, this metric provides the weakest independent construct validation among the three latent states.

For situational-awareness degradation, the reference is

$$z_{t,\mathrm{ref}}^{\mathrm{SA}} = \frac{(1 - P_t) + M_t}{2},$$

where P_t is probe accuracy and M_t is missed-hazard rate. Probe inaccuracy $(1 - P_t)$ contributes one-half of the reference score and provides a comparatively more distinct validation component, whereas missed-hazard rate contributes the other one-half and is behaviorally related to response and intervention patterns represented elsewhere in the framework. The RMSE decreases from 0.44 to 0.28:

$$R_{\mathrm{RMSE}}^{\mathrm{SA}} = \frac{0.44 - 0.28}{0.44} \times 100\% = 36.4\%,$$

with a corresponding squared-error reduction of

$$R_{\mathrm{MSE}}^{\mathrm{SA}} = \left[1 - \left(\frac{0.28}{0.44}\right)^2\right] \times 100\% \approx 59.5\%.$$

Using component errors e_P and e_M,

$$e_{\mathrm{SA}} = \frac{-e_P + e_M}{2},$$

so that

$$\mathrm{MSE}_{\mathrm{SA}} = \frac{1}{4}(\mathbb{E}[e_P^2] + \mathbb{E}[e_M^2] - 2\mathbb{E}[e_P e_M]).$$

Accordingly, the situational-awareness RMSE reduction has a comparatively stronger independent-validation component than the trust and risk metrics because probe accuracy is not itself an estimator input, although the missed-hazard component remains behaviorally coupled to the model. The equal 0.5 weights again describe the algebraic composition of the proxy and is not interpreted as the fraction of RMSE improvement attributable to independent versus overlapping information.

Taken together, the reductions from 0.49 to 0.31 for trust, 0.52 to 0.33 for risk sensitivity, and 0.44 to 0.28 for situational awareness correspond to approximately 36.7%, 36.5%, and 36.4% reductions in RMSE, respectively, or approximately 60.0%, 59.7%, and 59.5% reductions in squared error relative to Static LR. These values demonstrate substantially closer reconstruction of the predefined behavioral reference trajectories on held-out operators. The evidential strength of the RMSE reduction is lowest for risk sensitivity, for which both

https://doi.org/10.3390/computation14080192

reference components are coupled to modeled behavior or context; intermediate for trust miscalibration, for which confidence is directly shared with the estimator; and comparatively stronger for situational awareness, for which probe accuracy provides a more distinct reference component.

References

1. Sharma, G.; Jain, S.; Sharma, R.S. Path planning for fully autonomous uavs-a taxonomic review and future perspectives. *IEEE Access* **2025**, *13*, 13356–13379. [CrossRef]
2. Manju, S.; Prabha, M.; Farithkhan, A.; Sivagurunathan, G.; Michael Mahesh, K. Decentralized control design for UAV swarms communication. *Discov. Appl. Sci.* **2025**, *7*, 131. [CrossRef]
3. Yan, Y.; Li, X.; Shen, Z. A High-Precision Risk Analysis Model of Logistics UAV Network with Multiple Constraining Factors. *Trans. Nanjing Univ. Aeronaut. Astronaut.* **2024**, *41*, 218.
4. Zheng, Y.J.; Du, Y.C.; Sheng, W.G.; Ling, H.F. Collaborative human-UAV search and rescue for missing tourists in nature reserves. *Inf. J. Appl. Anal.* **2019**, *49*, 371–383. [CrossRef]
5. Perez, A.; Singh, A.; Mitchell, J.; Swadling, P. Strategies to Manage Human Factors in Mixed Reality Pilot Training: A Survey. *arXiv* **2025**, arXiv:2507.15526.
6. Perez, A.C.; Singh, A.K.; Swadling, P.; Mitchell, J. Strategies to Manage Human Factors in Mixed Reality Helicopter Pilot Training: A Systematic Literature Review. *Front. Virtual Real.* **2026**, *7*, 1815082. [CrossRef]
7. Hosseini, I.; Hossain, M.Z.; Zhang, Y.; Rahman, S. Deep learning model for simultaneous recognition of quantitative and qualitative emotion using visual and bio-sensing data. *Comput. Vis. Image Underst.* **2024**, *248*, 104121. [CrossRef]
8. Billing, E.; Fraboni, F.; Gualtieri, L.; Rosen, P.H.; Thorvald, P. Human factors and cognitive ergonomics in advanced industrial human-robot interaction. *Front. Robot. AI* **2025**, *12*, 1564948. [CrossRef] [PubMed]
9. Piessens, T.; van Paassen, M.M.; Mulder, M. UAV Haptic Interface for Dynamic Obstacle Avoidance. In Proceedings of the AIAA Scitech 2020 Forum, Orlando, FL, USA, 6–10 January 2020; p. 1112.
10. Lu, K.; Liu, W.; Svinin, M.; Wang, K.; Goncharenko, I. A System for UAV Obstacle Avoidance Control via Haptic Interface and Optical Tracking. In *Proceedings of the 2025 5th International Conference on Robotics, Automation, and Artificial Intelligence (RAAI), Singapore, 18–20 December 2025*; IEEE: Piscataway, NJ, USA, 2025; pp. 781–786.
11. Zak, Y.; Parmet, Y.; Oron-Gilad, T. Towards enhancement of unmanned aerial vehicle (UAV) operators' situation awareness: How often should their command and control map be updated. In Proceedings of the Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, Cambridge, UK, 23–26 March 2020; pp. 538–540.
12. Xiao, H.; Krisshna Kumar, P.; Pothuri, J.P.; Soni, P.; Butcher, E.; Chowdhury, S. An Open-source Hardware/Software Architecture and Supporting Simulation Environment to Perform Human FPV Flight Demonstrations for Unmanned Aerial Vehicle Autonomy. In Proceedings of the AIAA Aviation Forum and Ascend, Las Vegas, NA, USA, 29 July–2 August 2024; p. 4458.
13. Wu, C.; Guo, Z.; Zhang, J.; Mao, K.; Luo, D. Cooperative path planning for multiple UAVs based on APF B-RRT* algorithm. *Drones* **2025**, *9*, 177. [CrossRef]
14. Chen, J.; Zhang, X.; Tong, Y. *Cooperative Decision-Making Modeling for Manned and Unmanned Aerial Vehicles*; CRC Press: Boca Raton, FL, USA, 2026.
15. Sun, B.; Wu, J.; You, H.; Du, J. Multimodal feedback enhances human belief updating performance and reduces cognitive load in urban drone operation. *Front. Robot. AI* **2026**, *13*, 1707022. [CrossRef] [PubMed]
16. Cahillane, M.; Baber, C.; Morin, C. Human factors in UAV. In *Sense and Avoid in UAS: Research and Applications*; Wiley: Hoboken, NJ, USA, 2012; pp. 119–142.
17. Fan, J.; Chen, X.; Liang, X. UAV trajectory planning based on bi-directional APF-RRT* algorithm with goal-biased. *Expert Syst. Appl.* **2023**, *213*, 119137. [CrossRef]
18. Lai, C.C.; Yang, B.J.; Lin, C.J. Applying Reinforcement Learning for AMR's Docking and Obstacle Avoidance Behavior Control. *Appl. Sci.* **2025**, *15*, 3773. [CrossRef]
19. Stevšić, S.; Nägeli, T.; Alonso-Mora, J.; Hilliges, O. Sample efficient learning of path following and obstacle avoidance behavior for quadrotors. *IEEE Robot. Autom. Lett.* **2018**, *3*, 3852–3859. [CrossRef]
20. He, Y.; Cao, C.; Xiao, J. Day-ahead peak load probability density forecasting based on QRLSTM-DF considering exogenous factors. *IEEE Trans. Ind. Inform.* **2022**, *19*, 9447–9456. [CrossRef]
21. Li, F.; Li, Y.; Yuan, Y.; Li, J. UAV situation prediction technology based on Bayesian neural network in uncertain environment. In *Proceedings of the 2025 37th Chinese Control and Decision Conference (CCDC), Las Vegas, NA, USA, 21–25 July 2025*; IEEE: Boca Raton, FL, USA, 2025; pp. 5275–5280.
22. Galiautdinov, R. Information processing systems in UAV based on Bayesian filtering in conditions of uncertainty. *Int. J. Softw. Sci. Comput. Intell. (IJSSCI)* **2020**, *12*, 42–59. [CrossRef]

23. Aziz, I.; Zhang, R.; Ford, T.; Houtz, D.; Zhao, Y.; Watts, A.; Alipour, M. UAV-Based Remote Sensing of Soil Moisture Across Diverse Land Covers: Validation and Bayesian Uncertainty Characterization. *IEEE Trans. Geosci. Remote Sens.* **2025**, *63*, 1–18. [CrossRef]
24. Basnet, N.; Zahabi, M. Real-time cognitive workload assessment using non-intrusive methods: A systematic review. *Ergonomics* **2025**, 1–26. [CrossRef] [PubMed]
25. O'Hara, R.B.; Loftis, S.; Rando, C. Real-Time Biometric Monitoring for Cognitive Workload Detection: A Narrative Review of Applications in High-Demand Professions. *medRxiv* **2025**. [CrossRef]
26. Awumey, E.; Das, S.; Forlizzi, J. A systematic review of biometric monitoring in the workplace: Analyzing socio-technical harms in development, deployment and use. In Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, Rio de Janeiro, Brazil, 3–6 June 2024; pp. 920–932.
27. Faroni, M.; Beschi, M.; Pedrocchi, N. Safety-aware time-optimal motion planning with uncertain human state estimation. *IEEE Robot. Autom. Lett.* **2022**, *7*, 12219–12226. [CrossRef]
28. Liu, W.; Liang, X.; Zheng, M. Task-constrained motion planning considering uncertainty-informed human motion prediction for human-robot collaborative disassembly. *IEEE/ASME Trans. Mechatron.* **2023**, *28*, 2056–2063. [CrossRef]
29. Fisac, J.F.; Bajcsy, A.; Herbert, S.L.; Fridovich-Keil, D.; Wang, S.; Tomlin, C.J.; Dragan, A.D. Probabilistically safe robot planning with confidence-based human predictions. *arXiv* **2018**, arXiv:1806.00109.
30. Ahmadi, N.; Abbasian, P.; Hou, Q.; Powell, L.; Hammond, T. Physiological Sensing and Machine Learning Approaches Toward Anxiety Detection: A Systematic Review. *IEEE Sens. Rev.* **2025**, *3*, 89–116. [CrossRef]
31. Elcadi, G.H.; Lindskog, M.; Dahlkvist, R.; Lindberg, C.K.; Karlsson, T.; Lindqvist, J. Using heart rate variability and eye-tracking to differentiate simulator sickness severity rated by visual analogue scale. *Virtual Real.* **2026**, *30*, 112. [CrossRef]
32. Cao, Y.; Sweetser, P.; Chen, Z.; Zhu, X. Signals of Success and Struggle: Early Prediction and Physiological Signatures of Human Performance across Task Complexity. In Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems, Barcelona, Spain, 13–17 April 2026; pp. 1–23.
33. Peitzsch, A.G.; Geary, K.; Kong, Y.; Posada-Quintero, H.; Havard, D.; D'Angelo, W.R.; Chon, K.H. Multimodal Detection of Pain and Anticipation Anxiety from Ultra-Short Duration Wearable Sensors Measurements. *Sensors* **2026**, *26*, 3181. [CrossRef] [PubMed]
34. Ziegler, M.D.; Russell, B.A.; Kraft, A.E.; Krein, M.; Russo, J.; Casebeer, W.D. Computational models for near-real-time performance predictions based on physiological measures of workload. In *Neuroergonomics*; Academic Press: Cambridge, MA, USA, 2019; pp. 117–120.
35. Tarka, P. The comparison of estimation methods on the parameter estimates and fit indices in SEM model under 7-point Likert scale. *Arch. Data Sci.* **2017**, *2*, 1–6.
36. Dawes, J. Do data characteristics change according to the number of scale points used? An experiment using 5-point, 7-point and 10-point scales. *Int. J. Mark. Res.* **2008**, *50*, 61–104. [CrossRef]
37. Siegel, J.E. Data Proxies, the Cognitive Layer, and Application Locality: Enablers of Cloud-Connected Vehicles and Next-Generation Internet of Things. Doctoral Dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA, 2016.
38. Cepeda, C.; Tonet, R.; Osorio, D.N.; Silva, H.P.; Battegay, E.; Cheetham, M.; Gamboa, H. Latent: A flexible data collection tool to research human behavior in the context of web navigation. *IEEE Access* **2019**, *7*, 77659–77673. [CrossRef]
39. Dalilian, F. Empirical Investigation of Biometric and Behavioral Indicators of Human Performance in Bridge Inspection with Drone and Damage Detection with Artificial Intelligence Assistance. Doctoral Dissertation, The University of Iowa, Iowa City, IA, USA, 2024.
40. Dalilian, F.; Nembhard, D. Cognitive and behavioral markers for human detection error in AI-assisted bridge inspection. *Appl. Ergon.* **2024**, *121*, 104346. [CrossRef] [PubMed]
41. Li, Y. *Improving Inspector Training and Understanding Through Human-Centric Intelligence Methods for Drone-Assisted Bridge Inspection*; State University of New York: Stony Brook, NY, USA, 2023.
42. Kayayurt, B.; Yayla, I.; Yapici, A.; Küçükoğuz, C. Ground control station avionics software development in ANKA UAV. In *Proceedings of the 2011 IEEE/AIAA 30th Digital Avionics Systems Conference, Seattle, WA, USA, 16–20 October 2011*; IEEE: Piscataway, NJ, USA, 2011; p. 5B6–1.
43. Jovanović, M.; Starčević, D.; Jovanović, Z. Improving design of ground control station for unmanned aerial vehicle: Borrowing from design patterns. In *Proceedings of the 2010 36th EUROMICRO Conference on Software Engineering and Advanced Applications, Lille, France, 1–3 September 2010*; IEEE: Piscataway, NJ, USA, 2010; pp. 65–73.
44. Yang, D.; Dong, W.; Lu, W.; Liu, S.; Dong, Y. Unified monitor and controller synthesis for securing complex unmanned aircraft systems. *Drones* **2025**, *9*, 353. [CrossRef]
45. Pairet, È.; Hernández, J.D.; Carreras, M.; Petillot, Y.; Lahijanian, M. Online mapping and motion planning under uncertainty for safe navigation in unknown environments. *IEEE Trans. Autom. Sci. Eng.* **2021**, *19*, 3356–3378. [CrossRef]

46. Feng, J.; Zhang, J.; Zhang, G.; Xie, S.; Ding, Y.; Liu, Z. UAV dynamic path planning based on obstacle position prediction in an unknown environment. *IEEE Access* **2021**, *9*, 154679–154691. [CrossRef]
47. Wang, X.; Wang, X.; Zhou, Z.; Song, Y. Salient object detection with bayesian inference based on radar and camera fusion used in UAV obstacle avoidance. *Phys. Scr.* **2024**, *99*, 116010. [CrossRef]
48. De Rose, M. LiDAR-Based Dynamic Path Planning of a Mobile Robot Adopting a Costmap Layer Approach in ROS2. Doctoral Dissertation, Politecnico di Torino, Turin, Italy, 2021.
49. Behera, L.; Agarwal, S.; Sandhan, T.; Sharma, P.; Kumar, A.; Ranjan, A.; Watsa, S.; Singh, A.; Kasina, J.S. A cyber-physical system based unmanned ground vehicles for safety inspection and rescue support in an underground mine. *Int. J. Intell. Unmanned Syst.* **2025**, *13*, 92–128. [CrossRef]
50. Koeleman, A.S.M.; Ros, H.H.; Van Den Akker, T.J. Beat-to-beat interval measurement in the electrocardiogram. *Med. Biol. Eng. Comput.* **1985**, *23*, 213–219. [CrossRef] [PubMed]
51. Paiva, R.P.; Carvalho, P.; Couceiro, R.; Henriques, J.; Antunes, M.; Quintal, I.; Muehlsteff, J. Beat-to-beat systolic time-interval measurement from heart sounds and ECG. *Physiol. Meas.* **2012**, *33*, 177–194. [CrossRef] [PubMed]
52. Langley, P.; Bowers, E.J.; Murray, A. Principal component analysis as a tool for analyzing beat-to-beat changes in ECG features: Application to ECG-derived respiration. *IEEE Trans. Biomed. Eng.* **2009**, *57*, 821–829. [CrossRef] [PubMed]
53. Mosia, M. A Bayesian State-Space Approach to Dynamic Hierarchical Logistic Regression for Evolving Student Risk in Educational Analytics. *Data* **2025**, *10*, 23. [CrossRef]
54. Bollineni-Balabay, O.; Brakel, J.; Palm, F.; Boonstra, H.J. Multilevel hierarchical Bayesian versus state space approach in time series small area estimation: The Dutch travel survey. *J. R. Stat. Soc. Ser. A Stat. Soc.* **2017**, *180*, 1281–1308. [CrossRef]
55. Rivot, E.; Prévost, E.; Parent, E.; Baglinière, J.L. A Bayesian state-space modelling framework for fitting a salmon stage-structured population dynamic model to multiple time series of field data. *Ecol. Model.* **2004**, *179*, 463–485. [CrossRef]
56. Alghamdi, T.; Elgazzar, K.; Mostafi, S.; Ng, J. Improving spatiotemporal traffic prediction in adversary weather conditions using hierarchical bayesian state space modeling. In *Proceedings of the 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), Indianapolis, IN, USA, 19–22 September 2021*; IEEE: Piscataway, NJ, USA, 2021; pp. 3252–3258.
57. Hoffman, M.; Radul, A.; Sountsov, P. An adaptive-MCMC scheme for setting trajectory lengths in Hamiltonian Monte Carlo. In Proceedings of the International Conference on Artificial Intelligence and Statistics, Virtual, 13–15 April 2021; pp. 3907–3915.
58. Dombowsky, A. *Assessing the Quality of Posterior Samples from No-U-Turn Hamiltonian Monte Carlo*; McGill University: Montreal, QC, Canada, 2020.
59. Wang, Y.; Liu, Y. An adaptive probabilistic maintenance framework for decision planning optimization. In Proceedings of the AIAA Scitech Forum 2019, San Diego, CA, USA, 7–11 January 2019.
60. Quan, M. An advanced hybrid logistic regression model for static and dynamic mixed data classification. *IEEE Access* **2022**, *10*, 73623–73634. [CrossRef]
61. Rabiu, E.O.; Akande, D.O.; Adeyemo, Z.K.; Akanbi, I.A.; Obanisola, O.O. Three-state hidden Markov model for spectrum prediction in cognitive radio networks. *ABUAD J. Eng. Res. Dev. (AJERD)* **2024**, *7*, 425–435. [CrossRef]
62. Zhou, P.; Shi, W.; Tian, J.; Qi, Z.; Li, B.; Hao, H.; Xu, B. Attention-based bidirectional long short-term memory networks for relation classification. In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, Berlin, Germany, 7–12 August 2016; Volume 2: Short Papers, pp. 207–212.
63. Yang, J.; Huang, X.; Wu, H.; Yang, X. EEG-based emotion classification based on bidirectional long short-term memory network. *Procedia Comput. Sci.* **2020**, *174*, 491–504. [CrossRef]