

Bayesian Simulation Framework for Uncertainty-Aware Bird-Strike Risk Mapping

Indrajeet Aditya Roy
College of Engineering
Northeastern University
Boston, MA, United States
roy.i@northeastern.edu

Abstract—Bird strikes pose persistent safety and economic risks in aviation. We present a dynamic, uncertainty-aware framework for probabilistic bird-strike risk assessment in airport environments. The system combines (i) Kalman-filter state estimation for real-time bird tracking, informed by seasonal migration priors; (ii) Bayesian modeling of strike consequences via Beta-Binomial conjugacy on FAA Wildlife Strike data; and (iii) Gaussian-process spatial risk mapping that fuses real-time tracks with historical uncertainty to produce continuous risk surfaces with predictive variance. A sequential update scheme with temporal weighting adapts to evolving bird distributions while controlling the computational cost of GP refits. The framework yields path-wise risk metrics (means, credible bounds, exceedance probabilities) for operational corridors. In simulation, it captures diverging risk profiles across approach and departure paths while maintaining low tracking error, illustrating end-to-end uncertainty propagation and decision-relevant outputs.

I. INTRODUCTION

Bird strikes are a persistent and significant hazard in aviation [1], [2]. Approximately 90 % of reported bird strikes occur below 3,000 ft above ground level and within 9.3 km of an aerodrome, predominantly during aircraft take-off, climb-out, approach, and landing [3]. Effective mitigation strategies depend critically on accurate assessment of bird strike risk within the airport environment and along adjacent flight corridors.

Quantifying bird strike risk is challenging because of its spatiotemporal complexity and constantly changing conditions, as risk levels fluctuate with the geographic distribution of bird populations, species specific behaviors (e.g., flocking, altitude preferences), seasonal migration patterns, and weather conditions [6]. Accurately characterizing and modeling this evolving risk landscape is a challenge for aviation safety management systems. Most current bird strike risk tools handle uncertainty poorly as measurement errors, classification inaccuracies, and prediction uncertainties accumulate. They typically result in deterministic categorical outputs without uncertainty quantification [2]. Omitting uncertainty information can mislead time critical flight routing and safety decisions [1].

To address these limitations, this paper presents an integrated simulation framework for dynamic, probabilistic assessment of bird strike risk. The proposed approach utilizes external data sources to inform the underlying models: historical bird strike consequences are analyzed using the

U.S. Federal Aviation Administration (FAA) Wildlife Strike Database [3], which contains over 177,000 records detailing species, damage, and location; while seasonal bird movement patterns are informed by the *A global dataset of directional migration networks of migratory birds* dataset [7]. The framework integrates these data-derived insights with probabilistic state estimates derived from simulated real-time tracking data within a cohesive Bayesian workflow, combining multiple techniques to provide enhanced situational awareness and quantitative risk outputs. Key components and contributions include:

- **Probabilistic State Estimation:** Bird kinematic states (position, velocity) and their associated uncertainties (P_k) are estimated using Kalman filters (KF) that process noisy simulated radar observations. Incorporating seasonal bird movement behavioral priors derived from the bird migration networks dataset enhances predictive robustness and refines these estimates.
- **Probabilistic Consequence Modeling:** Bayesian parameter estimation via a Beta-Binomial conjugate model is utilized to analyze historical FAA bird strike data. This produces posterior Beta distributions for the probability of damage (p_{damage}) conditioned on context (e.g., species group, season), effectively quantifying the uncertainty ($\text{Var}(p_{\text{damage}})$) associated with historical risk factors rather than relying on deterministic frequencies.
- **Spatial Risk Mapping with Uncertainty Quantification:** Gaussian processes (GPs) perform nonparametric Bayesian regression to learn a continuous spatial risk map ($f(\mathbf{x})$). The GP is trained on input locations derived from Kalman filters (KF) position estimates ($\mathbf{x}_i$) paired with target risk scores (y_i) derived via Monte Carlo sampling from the posterior damage probability distributions (which capture historical uncertainty). To represent the positional uncertainty associated with the KF estimates, additional training points are placed in the proximity of position $\mathbf{x}_i$, guided by the KF error covariance matrix ($\mathbf{P}_k$), and assigned appropriately adjusted target values. This approach utilizes the flexibility of GPs to model complex spatial patterns based on uncertain inputs (y_i derived from sampling, and $\mathbf{x}_i$ represented with neighborhood points) and provides predictive uncertainty ($\sigma^2(\mathbf{x}^*)$)

across the airport, crucial for probabilistic assessment indicating multiple sources of uncertainty.

- **Dynamic Adaptation via Sequential Updates:** Recognizing that the airport environment is dynamic, the GP risk map is updated sequentially using a history buffer to manage computational load ($O(N^3)$ complexity) and temporal weighting heuristics to prioritize recent information. This involves systematically increasing the effective observation noise associated with older training data points (Δt_i) via an exponential decay mechanism during model refitting. By reducing the influence of outdated observations, this heuristic enables the spatial risk representation to adapt dynamically to changing simulated bird densities and movement patterns, simulating the variability of bird movement associated risk more realistically than static models.

II. METHODOLOGY

A. Iterative Simulation Workflow

The core simulation framework operates through an iterative workflow executed at each discrete time step, Δt . This sequential flow of data and processing within each time step allows the modeled spatial risk map to dynamically adapt to the evolving simulated bird distribution and tracking estimates.

- 1) **Ground-Truth State Evolution:** The simulation advances the ‘true’ state (position, velocity) for each bird. This step combines a kinematic prediction (based on previous velocity) with deterministic seasonal behavioral adjustments (e.g., migration bias derived from data detailed in Sec. II-B2), adds scaled (Δt) stochastic process noise representing unpredictable movement components, and enforces physical boundary constraints (e.g., minimum altitude) to generate the updated ground-truth positions for the current time step.
- 2) **Sensor Observation Simulation:** Based on the current ground-truth positions, the system simulates the observation process for the sensors. For each bird, its detectability by every sensor is evaluated. A bird is considered potentially detectable by a sensor only if its 3D distance is within the sensor’s specified maximum range. If within range, a probabilistic detection occurs based on a model where detection likelihood decreases quadratically with normalized distance ($P(\text{detect}) \approx 1 - (\text{distance}/\text{range})^2$). For each successful sensor detection, Gaussian noise, with standard deviation equal to the sensor’s specified accuracy is added to the bird’s true position to simulate measurement error. If a single bird is detected by multiple sensors in the same time step, these individual noisy measurements are combined by averaging their positions and associated accuracies to produce a single, consolidated observation point and effective accuracy for that bird.
- 3) **State Estimation (Filtering):** Each bird currently being tracked has its state updated by its dedicated Kalman

filter instance. The filter processes the combined sensor observation corresponding to that bird as input and then executes its Bayesian prediction-update cycle (detailed in Sec. II-B1) to optimally combine the prior prediction (informed by kinematics and bird movement behavioral priors) with the new measurement information, producing an updated posterior state estimate ($\mathbf{x}_k$: position, velocity) and its associated uncertainty covariance matrix ($\mathbf{P}_k$).

- 4) **Contextual Risk Potential Retrieval:** For risk assessment input, the system retrieves the pre-calculated probabilistic damage potential relevant to the current situation. Based on the tracked bird’s context (e.g., estimated species group) and the current simulation time (month/season), it looks up the corresponding posterior Beta distribution ($p_{\text{damage}} \sim \text{Beta}(\alpha_{\text{post}}, \beta_{\text{post}})$) derived from the historical FAA data analysis (detailed in Sec. II-C1).
- 5) **Spatial Risk Model Update:** The Gaussian Process model (Sec. II-D), which represents the continuous spatial risk is updated with the latest information. New training data points (Sec. II-D2) are constructed using the current KF state estimates ($\mathbf{x}_k$) as locations. The corresponding target risk values (y_k) are derived by performing Monte Carlo sampling (Sec. II-C2) on the retrieved Beta distributions (from step 4) to get an expected risk score. The GP model is then refit incorporating these new points and applying temporal weighting heuristics (Sec. II-D3) to adapt the risk map.
- 6) **Path Risk Evaluation:** The updated GP risk map is used to assess risk along predefined flight corridors. The GP’s posterior predictive distribution (Sec. II-D4) is queried at multiple points along each path to obtain location-specific probabilistic risk estimates (mean, variance, exceedance probabilities).

B. Bird Tracking and Behavior Modeling

Accurate estimation of bird positions and velocities in near real-time is crucial for dynamic risk assessment. This framework implements Kalman filters (KFs), a standard recursive Bayesian technique [8], [9], for this state estimation task, integrating kinematic modeling with bird movement behavioral priors derived from the bird migration networks dataset. Kalman filter was selected for its computational efficiency suitable for tracking multiple targets via recursive updates, theoretical optimality under linear/Gaussian assumptions, systematic Bayesian approach to combine model predictions with noisy sensor data, and ability to quantify state estimation uncertainty ($\mathbf{P}_k$), which is critical for informing subsequent probabilistic risk modeling.

- 1) **Kalman filter implementation:** Each detected bird is tracked using an independent Kalman filter instance, simplifying multi-target tracking. The Kalman filter estimates the bird’s 6D kinematic state vector, consisting of 3D position

(p_x, p_y, p_z) and 3D velocity (v_x, v_y, v_z) :

$$\mathbf{x} = [p_x \ p_y \ p_z \ v_x \ v_y \ v_z]^T \in \mathbb{R}^6 \quad (1)$$

Including velocity is essential for effective prediction. The Kalman filter's internal dynamics model implements a near-constant-velocity (NCV) kinematic model. The NCV model was selected due to its suitability for tracking targets based on sensor measurements that primarily provide position information (like radar), without directly measuring acceleration [10]. It provides a robust baseline for short term prediction, crucial for maintaining track continuity between potentially sparse sensor measurements, while relying on the process noise (Q) to accommodate deviations from constant-velocity motion. The Kalman filter operates recursively through prediction and update steps:

At each discrete time step interval (Δt), the Kalman filter first predicts the bird's next state based on the previous posterior estimate $\hat{\mathbf{x}}_{k-1}$ and the NCV motion model represented by the state transition matrix F . Applying this matrix produces the predicted state estimate before incorporating the latest measurement, denoted as $\mathbf{x}_k^-$:

$$\mathbf{x}_k^- = F \hat{\mathbf{x}}_{k-1} \quad (2)$$

Specifically, this state transition updates position based on the previously estimated velocity ($p_k^- = \hat{p}_{k-1} + \hat{v}_{k-1} \Delta t$) while assuming velocity remains constant ($v_k^- = \hat{v}_{k-1}$) during the interval. The uncertainty is propagated via the predicted state error covariance $\mathbf{P}_k^-$. This calculation incorporates the process noise covariance matrix $\mathbf{Q}$, which models the uncertainty introduced by random perturbations to the state dynamics, modeled as zero-mean Gaussian noise:

$$\mathbf{P}_k^- = F \mathbf{P}_{k-1} F^T + \mathbf{Q} \quad (3)$$

$\mathbf{Q}$ is constructed as a diagonal matrix whose elements represent the variances of the noise affecting each state component (position and velocity), scaled by the time step Δt . This structure assumes independent noise components and models the uncertainty originating from potential deviations from the constant velocity assumption (e.g., unexpected bird maneuvers/movements) during Δt .

When a new radar measurement $\mathbf{z}_k$ (position) is available, the Kalman filter corrects its prediction. It computes the innovation $\mathbf{y}_k$, the difference between the actual measurement and the measurement predicted from the prior state estimate ($H \cdot \mathbf{x}_k^-$):

$$\mathbf{y}_k = \mathbf{z}_k - H \cdot \mathbf{x}_k^- \quad (4)$$

where H is the observation matrix mapping state to measurement space. The optimal Kalman Gain K_k is calculated, balancing prediction uncertainty ($\mathbf{P}_k^-$) against measurement uncertainty (R , the measurement noise covariance). This gain determines how strongly the innovation corrects the predicted state, producing the updated posterior state estimate $\hat{\mathbf{x}}_k$ (the final estimate for time k) and its associated error covariance $\mathbf{P}_k$.

The posterior state estimate can be expressed explicitly with hat notation to denote the refined, post-update values:

$$\hat{\mathbf{x}}_k = [\hat{p}_x \ \hat{p}_y \ \hat{p}_z \ \hat{v}_x \ \hat{v}_y \ \hat{v}_z]^T \quad (5)$$

where the hat notation ($\hat{\cdot}$) distinguishes the posterior (after measurement update) from the predicted state. The covariance $\mathbf{P}_k$ quantifies the remaining uncertainty in this posterior estimate after incorporating the measurement information.

Spatial Anchors for GP Risk Mapping: The posterior mean position ($\hat{p}_x, \hat{p}_y, \hat{p}_z$) from the Kalman filter serves as a *spatial anchor* — the most probable 3D location of the bird at time step k . These spatial anchors form the critical bridge connecting state estimation to the Gaussian process risk model (Sec. II-D). Specifically, the posterior mean position provides the spatial coordinates $\mathbf{x}_i$ where the GP learns the risk distribution, while the covariance matrix $\mathbf{P}_k$ acts as an uncertainty prior that determines the confidence weight assigned to that observation. Regions where the Kalman filter exhibits high uncertainty (large $\mathbf{P}_k$) contribute less strongly to the learned spatial risk surface, ensuring consistency in uncertainty propagation across the entire Bayesian framework. Each Kalman-filtered bird position thus acts as a weighted probabilistic anchor in 3D space: highly certain estimates exert stronger influence on the GP risk field, while uncertain estimates are appropriately down-weighted.

2) Integration of Behavioral Priors: To improve predictive robustness and realistic bird movement behavior modeling beyond the NCV model, especially with sparse data, the KF's prediction step incorporates behavioral priors. These priors deterministically adjust the predicted velocity components in the state estimate $\mathbf{x}_k^-$ based on predefined seasonal tendencies and activity levels derived from ecological migration data [7], [10]. This produces more constrained and realistic trajectory estimates.

C. Historical Strike Consequence Analysis

To quantitatively characterize strike severity potential, historical data from the FAA Wildlife Strike Database [3] is utilized to estimate the conditional probability of aircraft damage given specific contextual factors, $P(\text{Damage}|\text{Strike}, \text{Context})$. A Beta-Binomial conjugate model was utilized for its ability to effectively quantify uncertainty in these probability estimates based on historical evidence.

1) Bayesian Damage Probability Estimation via Beta-Binomial Model: The damage probability (p_{damage}) conditional on various operational and environmental contexts is estimated. These contexts are defined by specific combinations of factors identified in the FAA bird strike data such as the bird's species group or family group, and the strike incident's month or season. Each historical strike event within a given context is modeled as a Bernoulli trial (damage or no damage). To represent uncertainty, our prior belief about p_{damage} for

any context is described by the conjugate Beta distribution, $\text{Beta}(\alpha_{\text{prior}}, \beta_{\text{prior}})$, whose pdf is:

$$\text{Beta}(p_{\text{damage}} | \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p_{\text{damage}}^{\alpha-1} (1 - p_{\text{damage}})^{\beta-1} \quad (6)$$

for $0 \leq p_{\text{damage}} \leq 1$, $\alpha, \beta > 0$. A weakly informative $\text{Beta}(1, 9)$ prior is applied across all contexts. This prior (mean = 0.1, equivalent to 10 prior observations: 1 damaging, 9 non-damaging) represents a mild baseline assumption of low damage probability, ensuring the posterior belief is primarily shaped by the context-specific FAA data evidence ($N_{\text{damage}}, N_{\text{no_damage}}$).

The likelihood of observing the historical data is modeled using the binomial distribution. Given N_{damage} damaging strikes and $N_{\text{no_damage}}$ non-damaging strikes within a specific context, where $N_{\text{total}} = N_{\text{damage}} + N_{\text{no_damage}}$, the binomial likelihood is:

$$P(N_{\text{damage}} | p_{\text{damage}}, N_{\text{total}}) = \binom{N_{\text{total}}}{N_{\text{damage}}} p_{\text{damage}}^{N_{\text{damage}}} (1 - p_{\text{damage}})^{N_{\text{no_damage}}} \quad (7)$$

where $\binom{N_{\text{total}}}{N_{\text{damage}}}$ is the binomial coefficient. This likelihood quantifies the probability of observing exactly N_{damage} damaging events out of N_{total} strikes, given a particular damage probability p_{damage} .

Utilizing the Beta-Binomial conjugacy, where the Beta prior is combined with the binomial likelihood, the posterior distribution for p_{damage} within this specific context remains a Beta distribution, $\text{Beta}(\alpha_{\text{posterior}}, \beta_{\text{posterior}})$, with updated parameters reflecting this context-specific evidence:

$$\begin{aligned} \alpha_{\text{posterior}} &= \alpha_{\text{prior}} + N_{\text{damage}} \\ \beta_{\text{posterior}} &= \beta_{\text{prior}} + N_{\text{no_damage}} \end{aligned} \quad (8)$$

This posterior $\text{Beta}(\alpha_{\text{posterior}}, \beta_{\text{posterior}})$ distribution represents the refined belief about p_{damage} under the specified conditions, integrating both prior knowledge and historical data relevant to that context. The variance of this posterior distribution, $\text{Var}(p_{\text{damage}})$, provides a direct quantification of the uncertainty in the p_{damage} estimate for that context. This uncertainty is inversely related to the amount of historical evidence ($N_{\text{damage}} + N_{\text{no_damage}}$) available for the context, decreasing as more data accumulates.

2) Combining Risk Factors and Propagating Uncertainty via Monte Carlo Sampling: The overall risk associated with a bird-path encounter depends on multiple factors: the inherent damage potential related to the bird's species and the time of year (represented by the uncertain posterior Beta distributions from Sec. II-C1), the geometry of the encounter (proximity), and a baseline historical risk rate. To integrate these diverse factors into a single risk score while effectively accounting for the uncertainty in the damage probabilities, a combined log-odds modeling and Monte Carlo sampling approach is utilized.

The different risk components are represented on a common log-odds scale to enable straightforward combination. This transformation facilitates an additive model where the distinct influences can be summed linearly. The baseline risk

and proximity effect contribute deterministic log-odds values, while the species and temporal (monthly or seasonal) effects contribute probabilistic log-odds values.

The log-odds transformation is employed because it enables all contributing factors to combine *additively* rather than *multiplicatively*. In probability space, effects such as species, proximity, and season would interact in complex, nonlinear ways that are difficult to interpret or separate. By mapping probabilities to log-odds, each factor contributes a linear adjustment to the total evidence for or against damage—conceptually analogous to the weights in logistic regression. This provides both mathematical convenience and interpretability: each term corresponds directly to an independent influence on the predicted log-odds of a damaging strike.

Formally, the overall risk function, denoted $f_{\text{risk}}(\cdot)$, combines both probabilistic and deterministic components. The total log-odds is expressed as:

$$L_{\text{total}} = L_{\text{baseline}} + \omega_{\text{species}} L_{\text{damage, species}} + \omega_{\text{temporal}} L_{\text{damage, temporal}} + L_{\text{proximity}} \quad (9)$$

where $L_{\text{damage, species}}$ and $L_{\text{damage, temporal}}$ represent the log-odds of damage derived from the respective posterior Beta distributions via the logit transformation:

$$L_{\text{damage}} = \log\left(\frac{p_{\text{damage}}}{1 - p_{\text{damage}}}\right) \quad (10)$$

More generally, this can be expressed as a summation over all contributing components:

$$L_{\text{total}} = \sum_i L_i \quad (11)$$

where each L_i represents a distinct factor (species, proximity, temporal, baseline) contributing additively in log-odds space. The final probability of a damaging strike is obtained by applying the inverse-logit (logistic) transformation:

$$f_{\text{risk}}(x) = \frac{1}{1 + e^{-L_{\text{total}}}} \quad (12)$$

This formulation enables the framework to integrate probabilistic insights from the Beta posterior over p_{damage} with deterministic contextual features. Because each term contributes additively in log-odds space, the influence of every component remains explicit (its numerical effect is directly represented in L_{total}), separable (each factor's contribution can be isolated and analyzed independently), and explainable (changes in risk can be traced back to specific environmental or operational drivers).

Critically, because p_{damage} is itself uncertain and modeled as $p_{\text{damage}} \sim \text{Beta}(\alpha_{\text{post}}, \beta_{\text{post}})$, the risk function becomes nonlinear in a random variable. The subsequent operations on p_{damage} —taking the logarithm (to obtain log-odds) and then applying the exponential (through the inverse-logit)—introduce nonlinear transformations. Hence, the expectation:

$$\mathbb{E}\left[\left(1 + e^{-L_{\text{total}}}\right)^{-1}\right] \quad (13)$$

has no closed-form solution, as the composition of log and exp functions destroys the conjugacy that made earlier

Bayesian updates tractable. Even though the Beta-Binomial stage was analytically solvable, this subsequent nonlinear mapping makes the overall risk integration analytically intractable.

Monte Carlo (MC) methods are therefore essential for this integration, allowing approximation of the distribution of the final risk score by repeatedly sampling from the input distributions (the Beta posteriors) and applying the risk calculation function to each set of samples. This process empirically reconstructs how the input uncertainties combine and influence the output uncertainty, bypassing the need for complex analytical derivations which may be intractable given the non-linear transformations (log-odds, inverse logit) involved. For each encounter, N_{MC} iterations are performed (we set $N_{MC} = 500$ in experiments). In each iteration j :

- 1) Sample Probabilities: Independent probability samples, $p_{\text{damage,species}}^{(j)}$ and $p_{\text{damage,temporal}}^{(j)}$, are drawn from the posterior Beta distributions reflecting the species/family and temporal (month/season) contexts.
- 2) Calculate Probabilistic Log-Odds Contributions: These sampled probabilities are converted into log-odds deviations from the prior baseline. For each sample j :

$$L_{\text{damage,species}}^{(j)} = \log\left(\frac{p_{\text{damage,species}}^{(j)}}{1 - p_{\text{damage,species}}^{(j)}}\right) - \log\left(\frac{0.1}{0.9}\right)$$

$$L_{\text{damage,temporal}}^{(j)} = \log\left(\frac{p_{\text{damage,temporal}}^{(j)}}{1 - p_{\text{damage,temporal}}^{(j)}}\right) - \log\left(\frac{0.1}{0.9}\right) \quad (14)$$

where the subtracted term $\log(0.1/0.9) \approx -2.197$ represents the log-odds of the Beta(1,9) prior mean (0.1). This centering ensures that species or temporal contexts with average damage rates contribute zero additional log-odds, while riskier contexts contribute positive values and safer contexts contribute negative values. These deviations are then weighted ($\omega_{\text{species}} \approx 1.0$, $\omega_{\text{temporal}} \approx 0.8$) based on empirically determined importance and literature support [17].

- 3) Add Deterministic Log-Odds Components: Two deterministic terms are added:
 - A fixed baseline log-odds component: $L_{\text{baseline}} = -2.5$. This empirically calibrated value corresponds to $p \approx 0.075$ (via $\text{inv_logit}(-2.5) = \frac{e^{-2.5}}{1 + e^{-2.5}}$), consistent with overall FAA damage rates (6-8%) [15]. Note that this differs slightly from the Beta(1,9) prior mean log-odds (-2.197) to better align with observed aggregate damage frequencies in the FAA dataset.
 - A proximity log-odds bonus computed as:

$$L_{\text{proximity}}^{(j)} = 2.0 \times (1 - d_{\text{norm}}^{(j)})^2 \quad (15)$$

where $d_{\text{norm}}^{(j)} \in [0, 1]$ is the normalized distance from the bird to the nearest flight path at iteration j . The quadratic term $(1 - d_{\text{norm}})^2$ ensures the risk increases sharply near flight paths (maximum of

2.0 at $d_{\text{norm}} = 0$) while decaying smoothly with distance, reflecting the heightened danger in aircraft operational areas [16].

- 4) Sum Total Log-Odds: The probabilistic and deterministic log-odds components are summed to produce the total log-odds for iteration j :

$$L_{\text{total}}^j = L_{\text{baseline}} + \omega_{\text{species}} L_{\text{damage,species}}^j + \omega_{\text{temporal}} L_{\text{damage,temporal}}^j + L_{\text{proximity}}^j \quad (16)$$

where the centered log-odds deviations from Step 2 are weighted and combined with the fixed baseline and dynamic proximity terms.

Repeating steps 1-4 N_{MC} times generates a distribution of $L_{\text{total}}^{(j)}$ samples, effectively representing the combined uncertainty. This distribution of total log-odds is then converted back to a distribution of risk probabilities, $P_{\text{risk}}^{(j)}$, using the inverse logit (logistic) function:

$$P_{\text{risk}}^{(j)} = \frac{1}{1 + \exp(-L_{\text{total}}^{(j)})} \quad (17)$$

The mean of this final probability distribution represents the expected risk score for the encounter, crucially incorporating the propagated uncertainty:

$$E[P_{\text{risk}}] = \frac{1}{N_{MC}} \sum_{j=1}^{N_{MC}} P_{\text{risk}}^{(j)} \quad (18)$$

This expected score serves as the target value for training the Gaussian Process spatial risk map:

$$y_k = E[P_{\text{risk}}] \quad (19)$$

where y_k is the risk score assigned to location $\mathbf{x}_k$ derived from the Kalman filter (Sec. II-D2).

D. Spatial Risk Mapping via Gaussian processes

Gaussian processes (GPs), a nonparametric Bayesian regression technique [?], are used to model the dynamic spatial risk. GPs define a probability distribution over possible risk functions $f(\mathbf{x})$, characterized by a mean and covariance (kernel) function, thus avoiding restrictive assumptions about the risk distribution's shape. The GP learns a continuous mapping from 3D spatial coordinates ($\mathbf{x} \in \mathbb{R}^3$) to the risk score $f(\mathbf{x})$, conditioned on training data consisting of Kalman filter position estimates paired with probabilistic risk scores derived from historical FAA bird strike data. Subsequently, for any query location $\mathbf{x}_*$, the GP yields a full posterior predictive distribution $p(f(\mathbf{x}_*)|\text{data})$, which provides both the mean predicted risk $\mu(\mathbf{x}_*)$ and the associated predictive uncertainty $\sigma^2(\mathbf{x}_*)$.

GPs were selected for this spatial risk modeling task due to several key advantages: (1) Flexibility: As nonparametric models, GPs can adaptively capture complex and potentially non-linear spatial risk distributions without strong prior assumptions about the risk function's form; (2) Uncertainty quantification: GPs provide predictive variance, $\sigma^2(\mathbf{x}_*)$, alongside the

mean, quantifying confidence based on factors such as training data density and proximity; and (3) Spatial interpolation: GPs enable continuous spatial risk estimation by utilizing kernel-based spatial correlations, allowing effective risk inference even from sparse bird detections.

1) GP Model Definition: To model the spatial risk field $f(\mathbf{x})$, the Gaussian Process prior employs a zero mean function and a carefully chosen covariance function (kernel), $k(\mathbf{x}_i, \mathbf{x}_j)$. The kernel defines prior assumptions about the risk function's spatial behavior. Given the complexity of bird strike risk – influenced by smooth environmental gradients, localized hotspots, and observation noise – a composite kernel is selected over simpler, stand-alone kernels. This allows us to separately model and control different aspects of the risk variation, leading to a more flexible and interpretable model. The composite kernel combines spatial structure and observation noise components:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \underbrace{\sigma_f^2 \cdot k_{\text{Matern}}(\mathbf{x}_i, \mathbf{x}_j \mid \boldsymbol{\ell}, \nu)}_{\text{Spatial Risk Structure}} + \underbrace{\sigma_n^2 \cdot \delta_{ij}}_{\text{Base Observation Noise}} \quad (20)$$

The core spatial structure of the risk field is modeled using a scaled Matern kernel ($\sigma_f^2 \cdot k_{\text{Matern}}$). The Matern kernel, with smoothness parameter $\nu = 1.5$, was chosen for its flexibility, allowing representation of continuous risk patterns that are potentially less smooth than assumed by alternatives like the RBF kernel, simulating realistic variations in factors like bird density near airport features. $\boldsymbol{\ell} = [\ell_x, \ell_y, \ell_z]$ enables the model to learn different spatial correlation ranges horizontally versus vertically, capturing the distinct structure of the airport runways and aircraft corridors environment. The overall expected variance of spatial risk fluctuations around the mean is controlled by the signal variance parameter σ_f^2 .

Variability in the training target risk scores (y_i) not captured by the smooth spatial structure is addressed by a White Kernel component ($\sigma_n^2 \cdot \delta_{ij}$), which models a baseline level of spatially independent observation noise. This noise component accounts for uncertainty and jitter originating from sources such as the Monte Carlo sampling process used to derive y_i (Sec. II-C2), underlying historical data stochasticity, and Kalman filter estimation errors associated with the input positions (Sec. II-B1). Additionally, to ensure the model adapts dynamically to changing bird activity over time, a temporal weighting heuristic (detailed in Sec. II-D3) is applied during model fitting. This heuristic effectively increases the observation noise variance attributed to older training points, reducing their influence and allowing the GP risk map to prioritize recent data reflecting current bird distributions and associated risks.

This composite kernel structure allows the GP to learn complex spatial dependencies using the flexible Matern kernel, account for observation noise in the training data via the White Kernel, and dynamically adapt to new information through the temporal weighting mechanism, providing an effective model for the time-varying spatial bird strike risk field.

2) Training Data Generation: The GP model is trained using data points $(\mathbf{x}_i, y_i)$ generated dynamically from the simulation outputs. For each relevant bird track, the input location $\mathbf{x}_i \in \mathbb{R}^3$ corresponds to the position estimated by its Kalman filter. The associated target value y_i represents a scaled estimate of the risk level at that location and time. This target y_i is derived by: (1) querying the pre-computed posterior Beta distributions (Sec. II-C1) relevant to the bird's estimated context (species group, family, time); (2) performing Monte Carlo sampling (Sec. II-C2) from these distributions to obtain an expected probabilistic risk score based on historical damage data; and (3) linearly scaling this probability score $p_i \in [0, 1]$ by a fixed constant factor ($S_{GP} \approx 2.5$) to obtain the final target value $y_i = p_i \cdot S_{GP}$ used for GP model training.

3) Sequential Updates and Temporal Adaptation: To enable the GP-based risk map to adapt to dynamic changes in bird distributions, the model is updated sequentially as new batches of training data (derived from recent bird tracks) become available. This avoids computationally prohibitive retraining on the entire observation history. To manage the $\mathcal{O}(N^3)$ complexity of GP fitting with N data points, a fixed-size history buffer stores the training data, with older or less relevant points being pruned when the buffer capacity is reached.

A temporal weighting heuristic is applied to prioritize recent observations. The effective observation noise variance associated with each training point $(\mathbf{x}_i, y_i)$ in the buffer is increased based on its age at each discrete time step Δt_i . This is achieved by adding an age-dependent term α_i (calculated using an exponential decay function $\alpha_i \propto 1 - \exp(-\lambda \Delta t_i)$, bounded between minimum and maximum values) to the diagonal of the kernel matrix during GP fitting. Consequently, older data points exert less influence on the posterior distribution compared to recent ones. This procedure is akin to discounting (power-prior) weighting of older observations. The GP is then refit using the current data buffer and these age-adjusted effective noise levels, producing an updated posterior risk map that indicates the most current understanding of the environment while systematically reducing the influence of past observations.

4) Probabilistic Prediction and Risk Assessment: Once trained, the GP provides predictions not just as single values, but as full probability distributions over the possible risk score f_* at any query location $\mathbf{x}_*$. Conditioned on the training data $\mathcal{D}$, this posterior predictive distribution is Gaussian, denoted $f_* | \mathbf{x}_*, \mathcal{D} \sim \mathcal{N}(\mu(\mathbf{x}_*), \sigma^2(\mathbf{x}_*))$. The mean $\mu(\mathbf{x}_*)$ represents the most likely risk score (point estimate), while the variance $\sigma^2(\mathbf{x}_*)$ quantifies the uncertainty associated with this prediction.

Utilizing this predictive distribution enables probabilistic assessments. For example, the probability that the risk f_* at $\mathbf{x}_*$ exceeds a threshold T can be computed using the predicted mean $\mu(\mathbf{x}_*)$ and standard deviation $\sigma(\mathbf{x}_*)$ with the standard

normal cdf (Φ):

$$P(f_* \geq T \mid \mathbf{x}_*, \mathcal{D}) = 1 - \Phi\left(\frac{T - \mu(\mathbf{x}_*)}{\sigma(\mathbf{x}_*)}\right) \quad (21)$$

E. Flight Path Risk Evaluation

To translate the spatial risk map into operationally relevant metrics, the framework evaluates the integrated probabilistic risk along predefined flight paths (aircraft approach/departure corridors). Each path is discretized into a sequence of query points $\mathbf{x}_{path}$.

At each query point, the current posterior predictive distribution from the trained GP model (Sec. II-D), $\mathcal{N}(\mu(\mathbf{x}_{path}), \sigma^2(\mathbf{x}_{path}))$, is obtained. From this distribution, point-wise metrics including the expected risk score $\mu(\mathbf{x}_{path})$, predictive uncertainty $\sigma(\mathbf{x}_{path})$, and the probabilities of exceeding risk thresholds ($P(f_* \geq T_{Medium})$, $P(f_* > T_{High})$) are calculated as described in Sec. II-D4.

Point-wise probabilistic assessments are aggregated along the path to derive metrics such as average expected risk score, maximum expected risk score, average predictive standard deviation, and average threshold exceedance probabilities, providing quantitative, uncertainty-aware information that can be utilized to compare the relative risk of different flight corridors under specific environmental conditions.


III. EVALUATION & RESULTS

The results discussed below (Fig. 1) originate from a specific Spring season scenario simulation. Key parameter for this simulation are summarized to provide context:

- **Environment, Flight Paths, and Duration:** The simulation was conducted within a 3D airport vicinity ($X, Y \in [-5, 5]$ km, Altitude $Z \in [0, 3]$ km) including four predefined flight corridors used for risk evaluation, depicted in risk maps (e.g., Fig. 2): Approach 09 (APP09, western approach, dashed blue line), Departure 27 (DEP27, eastern departure, solid green line), Approach 18 (APP18, southern approach, dashed blue line), and Departure 36 (DEP36, northern departure, solid green line). The simulation duration was 30 iterations, each representing a 10 s time step (Δt).
- **Initial State:** $N_{birds} = 15$ birds based on the FAA bird strike dataset were initialized including Goose (4), Duck (4), Small Bird (4), Gull (2), and Hawk (1), with randomized positions .
- **Bird Tracking:** Kalman filters estimated bird position and velocity using simulated observations from a 5-radar network (8-10 km range, 30-40 m accuracy). Tracking proved reliable, with average Kalman filter position error stabilizing below 10 m.
- **Probabilistic Risk Inputs:** The framework utilized historical FAA bird strike data analysis (Sec. II-C) and Monte Carlo sampling ($N_{MC} = 500$). Core risk parameters included: baseline log-odds=-2.5, proximity scale=+2.0, species scale=1.0, and temporal scale=0.8.

- **Spatial Risk Model (GP):** A Gaussian Process mapped risk across a 0.2km resolution grid. It utilizes a composite Matern kernel ($\nu = 1.5$, anisotropic lengths $\ell = [0.8, 0.8, 0.4]$ km) and a White kernel for noise. Target risk values for the GP were derived from scaled probabilities ($S_{GP} = 2.5$).
- **GP Adaptation Mechanism:** A critical setting was the slow temporal decay constant of 24 hours for the GP's sequential updates, managed with a 1000-point history buffer. This allowed the GP to maintain a longer memory of past risk patterns compared to faster decay settings.

A. Simulation Results: Spring Scenario



(a) Kalman filter position error vs. time step (m)



(b) Average Path Risk Probability ($P(\geq \text{Medium})$)

Fig. 1. Simulation results for the Spring scenario: (a) Kalman filter position estimation error over time, demonstrating tracking accuracy. (b) Evolution of average path risk probability ($P(\text{Risk} \geq \text{Medium})$) for approach and departure corridors, showing dynamic risk assessment.

B. Spatial Risk Map Evolution (Spring Scenario)

The simulation results for the Spring scenario highlight the framework's dynamic capabilities and the performance of its components. Figure 1a displays the Kalman filter (KF) position estimation error over the 30 simulation steps. After an initial convergence period within the first five steps, where the maximum error decreases sharply from over 55 m to below 20 m, the average estimation error stabilizes at a low level, consistently remaining below 10 m. The maximum error exhibits more variability, fluctuating mostly between 10-20 m, likely indicating instances of higher measurement noise from the simulated sensors. This demonstrates that the Kalman filters provided robust and accurate bird state estimates throughout the simulation, serving as a reliable input to the subsequent risk mapping stage.

Complementing the state estimation performance, Figure 1b shows the evolution of the average probability of encountering medium-or-high risk ($P(\text{Risk} \geq \text{Medium})$) along the four predefined flight paths, as estimated by the Gaussian Process (GP) risk map. A rapid increase occurs initially, with all paths reaching probabilities between 0.7 and 1.0 by step 5,




Fig. 2. Evolution of the GP spatial risk map ($P(\text{Risk} \geq \text{Medium})$) at altitude $z = 0.5$ km) during the Spring scenario at selected iterations. Estimated bird positions (black circles) and recent movement trails (yellow–purple dotted lines) influence the contours. Subcaptions report the mean predictive standard deviation over the slice.

indicating an immediate high-risk assessment based on early bird locations. Notably, the risk dynamics diverge significantly after approximately step 15. While the approach paths (APP09 and APP18) maintain very high average risk probabilities (generally >0.85), the departure paths exhibit a clear downward trend. Specifically, Departure 27 decreases to around 0.7, and Departure 36 shows a more pronounced decline, falling below 0.4 by the end of the simulation (step 30). This divergence, occurring despite the stable KF performance shown in Fig. 1a, indicates that the adaptive GP risk map, influenced by the slow (24 hour) temporal decay setting, successfully captured changes in the simulated spatial distribution of birds relative to different flight corridors over time. The framework effectively learned that the risk profile associated with departure paths diminished significantly in the latter half of this specific Spring scenario run, while approach paths remained high-risk areas.

The dynamic nature of the risk assessment is highlighted by the sequence of spatial risk map slices presented in Fig. 2. These maps show the evolution of $P(\text{Risk} \geq \text{Medium})$ at an altitude of $z = 0.5$ km for selected iterations (5, 10, 15, 20, 25, and 30) during the Spring simulation run. The underlying GP model's spatial risk distribution is directly influenced by the estimated bird positions (black circles with green 'x') and their recent movement history (yellow-to-purple dotted trails,

increasingly visible in Fig. 2d through 2f). The Mean Map Uncertainty (StdDev) value in each plot indicates the GP's overall predictive confidence across the displayed 2D slice. The GP predicts both a mean risk score $\mu(\mathbf{x}_*)$ and a predictive variance $\sigma^2(\mathbf{x}_*)$ (standard deviation $\sigma(\mathbf{x}_*)$) at each grid point $\mathbf{x}_*$; this value represents the average of the predictive standard deviations ($\text{mean}(\sigma(\mathbf{x}_*))$) calculated over all points in the specific map slice for that iteration. A lower value indicates higher average confidence (lower average uncertainty) in the GP's predictions across the spatial domain at that time step.

Consistent with the rapid initial increase observed in the path risk probabilities (Fig. 1b), the early risk maps (Fig. 2a and Fig. 2b) highlight formation of widespread high-risk areas (dark red, $P > 0.75$), predominantly along the aircraft approach and departure corridors. These initial hotspots correlate closely with the clusters of estimated bird positions near runway centerlines and path endpoints. During this phase, the Mean Map Uncertainty decreases substantially (from 0.369 at step 5 to 0.287 at step 10), indicating the GP quickly learning the initial high risk configuration based on these early observations.

The evolution of the risk maps (Fig. 2c through Fig. 2f) visually confirms the risk divergence previously identified in the path risk plot (Fig. 1b). While high risk ($P > 0.75$) generally persists along the aircraft approach corridors

(APP09 and APP18), corresponding to continued simulated bird presence near those areas, the risk intensity along the departure corridors (DEP27 and especially DEP36) noticeably diminishes over time. This can be inferred when comparing Fig. 2c to Fig. 2f, where the latter shows significantly lower risk probabilities (lighter orange/yellow) along the DEP36 corridor. This spatial differentiation demonstrates the GP effectively adapting to the simulated dynamics; as birds moved away from departure areas (visible from trails and changing positions), the temporal weighting heuristic reduced the influence of older data points, allowing the risk map to reflect the lower risk hazard in those zones based on more recent inputs. As the simulation progresses towards iteration 30 (Fig. 2f), while risk decreases in bird vacated areas, the remaining high-risk contours around the aircraft approach paths become spatially tighter and more concentrated compared to iteration 25 (Fig. 2e). This likely indicates the GP gaining confidence (Mean Map Uncertainty dropping from 0.247 to 0.207 between these steps) specifically around the current bird locations, leading to steeper risk gradients immediately surrounding these persistent clusters.

The sequence of risk maps (Fig. 2) demonstrate the Bayesian framework's effectiveness. It effectively translates the discrete and uncertain bird position estimates derived from the Kalman filters into a continuous spatial risk field where predictive uncertainty is quantified via the Gaussian process. These dynamic maps provide spatial context for interpreting the aggregated path risk metrics (Fig. 1b), visualizing how the framework captures the time-varying nature of simulated bird distributions and the resulting concentration of risk around persistent bird locations within critical airport airspace environment.

REFERENCES

- [1] National Business Aviation Association (NBAA), "Wildlife Strike Response," <https://nbaa.org/aircraft-operations/safety/in-flight-safety/wildlife-strike-response/>.
- [2] Aircraft Owners and Pilots Association (AOPA), "Bird and wildlife strikes," <https://www.aopa.org/training-and-safety/active-pilots/safety-and-technique/bird-and-wildlife-strikes>.
- [3] U.S. Federal Aviation Administration (FAA), "Wildlife Strike Database," <https://wildlife.faa.gov/>.
- [4] J. Thorpe, "Fatalities and destroyed civil aircraft due to bird strikes, 1912-2002," in *Proc. Int. Bird Strike Comm. Meeting*, vol. 5, pp. 1-1, 2003.
- [5] J. R. Allan, "The costs of bird strikes and bird strike prevention," in *Human-wildlife conflicts: economic considerations*, Filion, FL; R. F. Association Eds., 2002.
- [6] C. Soldatini, Y. V. Albores-Barajas, R. P. Soriano-Redondo, and N. Saino, "A probabilistic approach to assess birdstrike risk in Italian airports," *Eur. J. Wildl. Res.*, vol. 57, pp. 99-108, 2011.
- [7] Y. Qiu, "A global dataset of directional migration networks of migratory birds," figshare, Dataset, 2024. [Online]. Available: <https://doi.org/10.6084/m9.figshare.26162269.v3>
- [8] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Trans. ASME J. Basic Eng.*, vol. 82, no. 1, pp. 35-45, Mar. 1960.
- [9] G. Welch and G. Bishop, "An introduction to the Kalman filter," University of North Carolina at Chapel Hill, Dept. of Computer Science, TR 95-041, 1995. [Online]. Available: https://www.cs.unc.edu/~welch/media/pdf/kalman_intro.pdf

- [10] W. D. Blair and Y. Bar-Shalom, "On the NCA Versus NCV Models in Tracking Maneuvering Targets," 2023 IEEE Radar Conference (Radar-Conf23), San Antonio, TX, USA, 2023, pp. 1-6, doi: 10.1109/Radar-Conf2351548.2023.10149710.
- [11] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for hyperparameter optimization," in *Adv. Neural Inf. Process. Syst. 24 (NIPS 2011)*, pp. 2546-2554, 2011.
- [12] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Adv. Neural Inf. Process. Syst. 30 (NIPS 2017)*, pp. 4765-4774, 2017.
- [13] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [14] D. W. Hosmer, S. Lemeshow, and R. X. Sturdivant, *Applied Logistic Regression*, 3rd ed. Hoboken, NJ, USA: Wiley, 2013.
- [15] R. A. Dolbeer, M. J. Begier, J. L. Seubert, S. E. Wright, and R. B. Walz, "Wildlife strikes to civil aircraft in the United States, 1990-2023," Federal Aviation Administration, Office of Airport Safety & Standards and USDA Wildlife Services, Serial Rep. 30, 2024. [Online]. Available: https://www.faa.gov/airports/airport_safety/wildlife/wildlife-strike-report-1990-2023-USDA-FAA
- [16] I. C. Metz, J. Ellerbroek, T. Mühlhausen, D. Kügler, S. Kern, and J. M. Hoekstra, "The efficacy of operational bird strike prevention," *Aerospace*, vol. 8, no. 1, art. 17, Jan. 2021, doi: 10.3390/aerospace8010017.
- [17] C. Nilsson, F. A. La Sorte, A. M. Dokter, K. G. Horton, B. M. Van Doren, J. J. Kolodzinski, J. Shamoun-Baranes, and A. Farnsworth, "Bird strikes at commercial airports explained by citizen science and weather radar data," *J. Appl. Ecol.*, vol. 58, no. 10, pp. 2029-2039, 2021, doi: 10.1111/1365-2664.13971.