

Generative AI Against Poaching: Latent Composite Flow Matching for Poaching Prediction

Lingkai Kong^{1*}, Haichuan Wang^{1*}, Charles A. Emogor^{1,2}, Vincent Börsch-Supan¹, Lily Xu³, Milind Tambe¹

¹School of Engineering and Applied Sciences, Harvard University

²Department of Computer Science and Technology, University of Cambridge

³Department of Industrial Engineering and Operations Research, Columbia University

Abstract

Poaching poses significant threats to biodiversity. A valuable step in reducing poaching is to forecast poacher behavior, which can inform patrol deployment and other conservation interventions. Existing poaching prediction methods based on linear models or decision trees lack the expressivity to capture complex, nonlinear spatiotemporal patterns. Recent advances in generative modeling, particularly flow matching, offer a more flexible alternative. However, training such models on real-world poaching data faces two central obstacles: imperfect detection of poaching events and limited data. To address imperfect detection, we integrate flow matching with an occupancy-based detection model and train the flow in latent space to infer the underlying occupancy state. To mitigate data scarcity, we adopt a composite flow initialized from a linear-model prediction rather than random noise which is the standard in diffusion models, injecting prior knowledge and improving generalization. Evaluations on datasets from two national parks in Uganda show consistent gains in predictive accuracy.

Extended version — <https://arxiv.org/abs/2508.14342>

1 Introduction

Poaching (i.e., illegal hunting) is a major direct driver of global biodiversity loss. It affects a wide range of species and undermines both ecological functioning (Abernethy et al. 2013) and, over time, the livelihoods of local communities (Coad et al. 2019). Multiple reports highlight the scale of these losses. Across Africa, about 100,000 elephants were illegally killed from 2010–2012 for ivory (Wittemyer et al. 2014). Pangolins are targeted for their scales and meat, with about 8.5 million taken for illegal trade in West and Central Africa from 2014–2021 (Emogor et al. 2021; United Nations Office on Drugs and Crime 2020). Predators are also heavily affected: between 2000 and 2022, at least 3,377 tigers were seized globally for their skins and other body parts (Sohl 2022). Taken together, these figures underscore the need to direct limited patrol resources to the places where they can best prevent poaching, in line with global reports on

*These authors contributed equally.



Figure 1: Well-hidden snares and rangers conducting a patrol to locate them. Photos credit: Uganda Wildlife Authority

wildlife trafficking (United Nations Office on Drugs and Crime 2024).

To strengthen protection under such severe losses, an effective strategy is to forecast poacher behavior and use these predictions to guide patrol planning. Machine learning models have been applied to estimate poaching risk and support data-driven patrol allocation. However, existing approaches typically rely on classical models—such as linear models (Xu et al. 2021) and decision trees (Thiault et al. 2020; Kar et al. 2017)—which are unable to capture the complex spatial patterns and the strategic adaptations of poachers in response to patrols. These models often make simplifying assumptions about spatial independence, which can lead to inaccurate risk forecasts and, in turn, less effective patrol deployment.

Recent advances in generative AI, particularly flow matching (Lipman et al. 2023) and diffusion models (Ho, Jain, and Abbeel 2020), offer powerful distribution-fitting capabilities that can model complex, high-dimensional patterns. These models are well-suited to capturing the nonlinear, strategic, and spatially correlated nature of poaching behavior. However, applying them to real-world poaching datasets presents two key challenges. First, the amount of historical data is typically small, especially in parks with limited monitoring capacity, whereas such generative models are usually data-hungry. Second, detection of poaching events is imperfect: many illegal activities go unnoticed during patrols, and negative labels indicating the absence of poaching are unreliable due to the difficulty of detecting well-hidden signs such as snares in dense vegetation (Keane, Jones, and Milner-Gulland 2011).

To address these challenges, we propose WILDFLOW, a framework based on *latent composite flow matching*. Building on occupancy models widely used in ecology (MacKenzie et al. 2002), we represent poaching activity as a latent state, distinct from the probability of detection during patrols. This separation enables explicit handling of imperfect detection: the latent state evolves through flow matching, while detection is modeled as the observation process in an occupancy framework. To mitigate data scarcity, we further design a composite flow that initializes the generative process from a linear model’s prediction rather than random noise, which is standard in diffusion models. By embedding domain knowledge in the initialization, our method improves generalization and supports effective learning in small-data settings.

We summarize our contributions as follows: (1) We present WILDFLOW, the first application of generative AI to poaching prediction. (2) We propose a latent composite flow matching framework that jointly addresses the challenges of limited data and imperfect detection in conservation. (3) On real-world data from Murchison Falls and Queen Elizabeth National Parks in Uganda, our method improves the area under the precision–recall curve (AUPR) by 7.0% and 10.2%, respectively. We also outline steps toward responsible deployment, including field pilots, and careful rollout.

2 Related Works

AI for poaching prediction. Yang et al. (2014) introduced the Protection Assistant for Wildlife Security (PAWS) using a subjective-utility quantal response (SUQR) behavioral model. To explicitly account for imperfect detection, Nguyen et al. (2016) proposed a two-layer Bayesian network with latent variables. However, Kar et al. (2017) later found this approach brittle in practice due to model complexity; an ensemble of decision trees achieved higher accuracy with lower runtime and was validated in a one-month field test. To combine the strengths of prior methods, Gholami et al. (2017) proposed a geo-clustering technique that yields a hybrid of multiple Markov random fields with a bagging ensemble of decision trees, supported by a five-month controlled field test. To further address nonuniform uncertainty stemming from uneven patrol effort, Gholami et al. (2018); Xu et al. (2020) trained ensembles of weak learners—decision trees and Gaussian processes—on data stratified by patrol intensity, with deployment via the SMART conservation platform (SMART 2013). Finally, Xu et al. (2021) used logistic regression to quantify the deterrence effects of ranger patrols on poaching risk.

Despite these practical successes, most approaches still rely on classical models with limited capacity to capture high-dimensional, nonlinear spatial patterns; few make use of modern deep learning. Most recently, Kong et al. (2025) applied diffusion models to green security problems. Our work differs in three key ways: (i) Their focus is on robust patrol optimization for general green security, whereas we focus specifically on poaching risk forecasting for conservation. (ii) They do not account for imperfect detection, which is a critical challenge in conservation settings, while we explicitly model the visit-level detection process. (iii) Their

features are restricted to historical patrol effort, whereas we incorporate rich environmental covariates. In this richer feature space, diffusion models struggle due to their data-hungry nature. To address this, we propose a composite flow model that improves data efficiency.

Generative AI for wildlife conservation In camera-trap vision, neural generative models have been applied directly to wildlife datasets: CycleGAN variants translate between sensor domains (e.g., visible and near-infrared) and augment rare species to improve few-shot classification (Gao et al. 2020; Zhang et al. 2023). Beyond imagery, generative audio models have been developed to synthesize wildlife vocalizations and strengthen bioacoustic monitoring under data scarcity. Early work explored class-conditional GANs for animal-sound augmentation (Kim et al. 2023), followed by diffusion-based pipelines that generate birdsong spectrograms to improve classifier accuracy (Gibbons et al. 2024), and more recently, diffusion models that synthesize anuran calls (Manrique, Gómez, and Hernández-Romero 2025). Large language models (LLMs) have also been applied to detect illegal wildlife trafficking on online marketplaces by generating pseudo-labels for unlabeled listings (Barbosa et al. 2025). In contrast, our work focuses on generative modeling for poaching prediction in protected areas.

3 Task, Challenges, and Preliminaries

3.1 Task Definition

We study Murchison Falls National Park (MFNP, $\sim 5,000 \text{ km}^2$) and Queen Elizabeth National Park (QENP, $\sim 2,500 \text{ km}^2$) in Uganda. Both parks are critical for ecotourism and conservation and provide habitat for elephants, giraffes, hippos, and lions (Critchlow et al. 2015).

We partition each park into a grid of $1 \times 1 \text{ km}$ cells, indexed by $i \in \{1, \dots, N\}$, and time into monthly intervals $t \in \{1, \dots, T\}$. For each cell–month, we construct a feature vector $\mathbf{x}_{i,t} \in \mathbb{R}^d$ that includes *static* geospatial attributes (aspect, elevation, surface water, seasonal water cover, and distances to rivers, patrol posts, and towns/villages) as well as *dynamic* covariates (monthly average precipitation, temperature, and net primary productivity). The latent occupancy state, i.e., whether poaching activity occurs, is denoted by $z_{i,t} \in \{0, 1\}$.

To combat poaching, rangers conduct GPS-tracked patrols, record observations, and remove wire snares that would otherwise trap wildlife. Within month t , cell i may be visited $J_{i,t}$ times. For visit $j \in \{1, \dots, J_{i,t}\}$, let $a_{i,t,j} \geq 0$ denote the *visit-level patrol effort* (kilometers patrolled), and let $y_{i,t,j} \in \{0, 1\}$ indicate whether illegal activity was detected during that visit.

Objective. Our goal is to model and forecast the true poaching risk, $p(z_{i,t} = 1)$, using environmental covariates $\{\mathbf{x}_{i,t}\}$ and historical patrol effort $\{a_{i,t-1,j}\}$, while explicitly accounting for the current month’s visit-level effort $\{a_{i,t,j}\}$.

3.2 Challenges

Complex poacher behavior Poaching dynamics emerge from intertwined social and ecological drivers and adapt

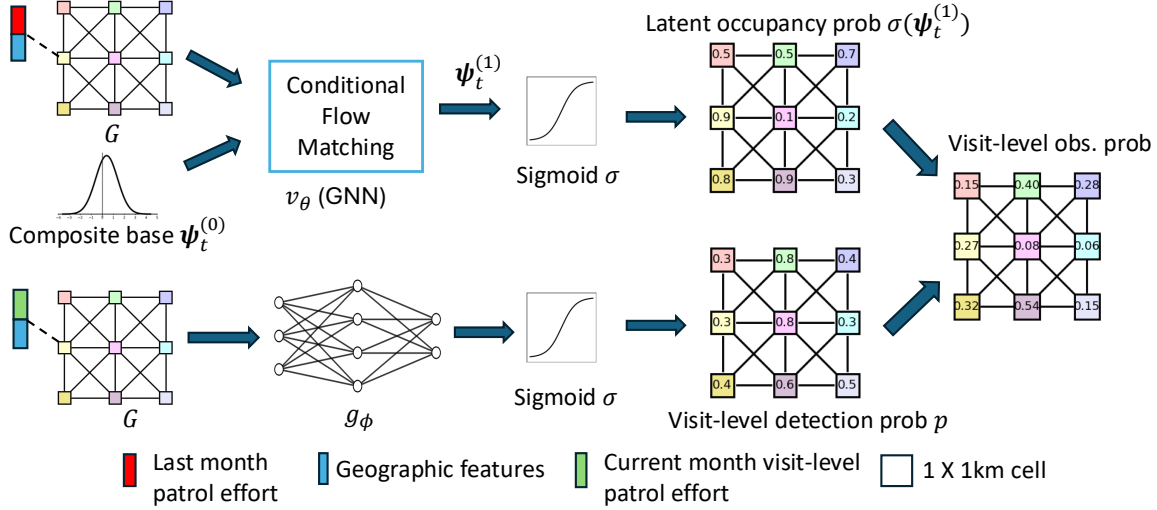


Figure 2: Overview of the latent composite flow framework. **Upper branch (latent generation):** 1×1 km cells are nodes in a graph G with edges between adjacent cells. For month t , node features comprise geospatial covariates $\mathbf{x}_{i,t}$ and last month’s aggregated patrol effort $a_{i,t-1}^m = \sum_j a_{i,t-1,j}$. A *composite base* initializes the latent logits $\psi_t^{(0)}$ from a pretrained linear occupancy model with Gaussian perturbation. A graph-conditional velocity field v_θ transports $\psi_t^{(0)}$ to $\psi_t^{(1)}$ via conditional flow matching; $\sigma(\psi_t^{(1)})$ yields occupancy probabilities. **Lower branch (detection):** the visit-level detector $g_\phi(\mathbf{x}_{i,t}, a_{i,t,j})$ uses geospatial features and current-month visit effort to produce detection probabilities $\{p_{i,t,j}\}$, which are combined with occupancy in the occupancy–detection likelihood.

strategically to enforcement. Prior work highlights a deterrence–displacement effect: stronger patrol presence lowers activity within a cell but shifts risk to neighboring cells, leading to spatial spillovers (Dancer et al. 2022; Rytwinski et al. 2024; Xu et al. 2021). These dynamics yield non-stationary and spatially coupled patterns that simple discriminative models struggle to capture.

Imperfect detection Observed incidents reflect *detection*, not true occurrence. Because snares and traps are often well hidden, the likelihood of detection depends on factors such as habitat, visibility, and patrol effort (Dobson et al. 2020). As a result, the data contain many false negatives and are biased by patrol routing. Models that overlook this distinction conflate absence with non-detection, leading to biased risk estimates. It is therefore essential to separate true poaching occurrence from the detection process.

Small data Although some parks have up to 20 years of patrol records, strong nonstationarity limits the *effective* sample size: staffing levels, prices for illegal wildlife products and land–use change alter incentives and accessibility, making older data weak predictors of current risk. Following prior work (Xu et al. 2020), we therefore restrict training to the most recent ≤ 3 years to mitigate temporal drift. In addition, poaching patterns, ranger resources, and landscape features differ markedly across parks; naive cross-park pooling induces severe dataset shift, so data cannot simply be combined to train a single universal model.

3.3 Preliminaries on Flow Matching

To capture complex spatial structure, we adopt *flow matching* (FM), a generative framework that learns a time-dependent velocity field to transport a base distribution toward the data distribution (Lipman et al. 2023). FM enables faster sampling than diffusion models (Ho, Jain, and Abbeel 2020), while often achieving comparable or improved quality.

In this work, we focus on the *conditional* version. At flow time $s=0^1$, FM draws a base sample $\psi^{(0)} \sim p_0(\cdot | \mathbf{c})$ and evolves it under a conditional velocity field $\frac{d\psi^{(s)}}{ds} = v_\theta(\psi^{(s)}, s; \mathbf{c})$, where \mathbf{c} denotes contextual information. Integrating from $s=0$ to $s=1$ yields $\psi^{(1)}$, a sample from $p_{\text{data}}(\cdot | \mathbf{c})$.

A key distinction from diffusion models is the choice of the initial distribution: FM permits p_0 to be any distribution, whereas diffusion models must begin with a standard Gaussian. Training FM commonly involves minimizing the loss

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{\substack{\mathbf{c} \sim p_{\text{data}}(\mathbf{c}), \psi^{(1)} \sim p_{\text{data}}(\cdot | \mathbf{c}), \\ \psi^{(0)} \sim p_0(\cdot | \mathbf{c}), s \sim \mathcal{U}(0,1)}} \left\| v_\theta(\psi^{(s)}, s; \mathbf{c}) - (\psi^{(1)} - \psi^{(0)}) \right\|_2^2, \quad (1)$$

where $\psi^{(s)} = (1-s)\psi^{(0)} + s\psi^{(1)}$. The linear interpolation ensures smooth trajectories, reducing ODE discretization error and improving sampling efficiency.

While FM can capture complex spatial patterns, it does not by itself address (i) imperfect detection or (ii) small data.

¹Throughout, t indexes months in the poaching data, while $s \in [0, 1]$ parameterizes the generative flow.

4 Proposed Method

In this section, we propose WILDFLOW to address the challenges of imperfect detection and limited data.

To handle imperfect detection, our framework builds on ecological occupancy models (MacKenzie et al. 2002), which treat the true presence of poaching as a hidden state and explicitly model the probability of detection during patrols. Leveraging flow matching’s ability to capture complex, high-dimensional spatial patterns, we model the generative process of the latent occupancy state (i.e., the true presence of poaching) and add an explicit visit-level detection model on top. This design allows us to distinguish true absence from non-detection, thereby reducing dataset bias.

To address the limited-data challenge, we warm-start flow matching from a composite base initialized by a simple linear model. This initialization embeds domain knowledge, improves generalization, and enables more data-efficient learning than starting from random noise.

4.1 Data Generation Process

Figure 2 illustrates the data generation process of our framework. The design has two components: the *upper branch*, which generates the probability of the latent occupancy states (i.e., the true but unobserved presence of poaching), and the *lower branch*, which models the detection process during ranger patrols.

Flow matching for latent occupancy generation. We represent each park as a grid graph $G = (V, E)$ of 1×1 km cells, where nodes correspond to cells and edges connect neighbors. For month t , each cell $i \in V$ has features $\mathbf{c}_{i,t} = [\mathbf{x}_{i,t}; a_{i,t-1}^m]$, combining geographic covariates $\mathbf{x}_{i,t}$ and the previous month’s aggregated patrol effort $a_{i,t-1}^m = \sum_j a_{i,t-1,j}^m$. This lagged effort captures patrol deterrence effects.

We place a conditional flow prior on the *joint* vector of occupancy logits $\boldsymbol{\psi}_t^{(1)} \in \mathbb{R}^N$ ($N=|V|$) and evolve it via

$$\frac{d\boldsymbol{\psi}_t^{(s)}}{ds} = v_\theta(\boldsymbol{\psi}, s; G, \mathbf{C}_t), \quad \boldsymbol{\psi}_t^{(0)} \sim p_0(\cdot | \mathbf{C}_t), \quad s \in [0, 1],$$

where \mathbf{C}_t is the collection of $\mathbf{c}_{i,t}$ over all nodes and v_θ is a message-passing graph neural network (GNN) (Kipf and Welling 2017). Message passing propagates information across neighboring cells during the flow, enabling the model to capture spatial spillovers and local correlations in poaching risk. The final logits are mapped through a sigmoid to yield latent occupancy probabilities $r_{i,t} = \sigma(\boldsymbol{\psi}_{i,t}^{(1)})$.

Warm-starting from a composite base distribution. Instead of initializing the flow from uninformative Gaussian noise, we warm-start from a *composite base*:

$$\boldsymbol{\psi}_t^{(0)} = b_\eta(\mathbf{C}_t) + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \mathbf{I}),$$

where b_η is a lightweight pretrained linear occupancy predictor and \mathbf{C}_t collects node features. This base injects domain knowledge from simple ecological models while still allowing variability through the Gaussian perturbation, improving data efficiency in small-sample settings.

Visit-level detection model. The lower branch of Fig. 2 describes how rangers detect snares during patrols. Within each month t , a cell i may be visited several times, where each visit j involves patrol effort $a_{i,t,j}$ and results in a binary outcome $y_{i,t,j} \in \{0, 1\}$ indicating whether snares were found. We model the detection probability with a lightweight neural head g_ϕ that takes environmental features and patrol effort as inputs: $p_{i,t,j} = \sigma(g_\phi(\mathbf{x}_{i,t}, a_{i,t,j}))$. Since detections are only possible when the site is truly occupied, the model naturally separates *true absence* from *non-detection*. This design reduces bias in the observations and is consistent with ecological occupancy modeling practices.

Let $S_{i,t} = \mathbf{1}\{\exists j : y_{i,t,j} = 1\}$ indicate that *at least one* detection occurred in cell i during month t . Marginalizing $z_{i,t}$ gives the visit-level occupancy–detection log-likelihood:

$$\begin{aligned} \mathcal{L}_{\text{occ}}(i, t) = & \log \left[(1 - r_{i,t})(1 - S_{i,t}) \right. \\ & \left. + r_{i,t} \prod_{j=1}^{J_{i,t}} p_{i,t,j}^{y_{i,t,j}} (1 - p_{i,t,j})^{1 - y_{i,t,j}} \right]. \end{aligned}$$

We provide the full derivation of \mathcal{L}_{occ} in Appendix A.1 of the extended version.

4.2 Training and Inference Algorithm

As discussed in Section 3.3, flow matching (FM) requires samples from the *target* distribution. Since the true latent logits $\boldsymbol{\psi}_t^{(1)}$ are unobserved, direct training is infeasible. Inspired by stable diffusion (Rombach et al. 2022), we adopt a two-stage procedure.

Stage 1: Encoder–detector training. In the first stage, we jointly train two components: (i) an encoder f_ω that maps monthly observations, environmental features and current month patrol effort into latent occupancy logits $\hat{\boldsymbol{\psi}}_t$, and (ii) a detection head g_ϕ that predicts visit-level detection probabilities. The parameters (ω, ϕ) are optimized by maximizing the summed occupancy–detection likelihood $\sum_{i,t} \mathcal{L}_{\text{occ}}(i, t)$.

The resulting encoder produces surrogate logits $\hat{\boldsymbol{\psi}}_t^{(1)}$, which serve as training targets for the flow model in Stage 2.

Stage 2: Latent flow training. In the second stage, we freeze the encoder–detector pair (ω, ϕ) and train a conditional flow model v_θ to transport samples from the composite base $\boldsymbol{\psi}_t^{(0)}$ toward the surrogate logits $\hat{\boldsymbol{\psi}}_t^{(1)}$.

Following the linear–path flow matching objective in Eq. (1), we sample $s \sim \mathcal{U}[0, 1]$ and form an interpolated state: $\boldsymbol{\psi}_t^{(s)} = (1 - s)\boldsymbol{\psi}_t^{(0)} + s\hat{\boldsymbol{\psi}}_t^{(1)}$. The velocity field v_θ is then trained to predict the straight-line displacement: $\|v_\theta(\boldsymbol{\psi}_t^{(s)}, s; G, \mathbf{C}_t) - (\hat{\boldsymbol{\psi}}_t^{(1)} - \boldsymbol{\psi}_t^{(0)})\|_2^2$. Here (G, \mathbf{C}_t) provide conditioning context from the graph and features.

Inference. For a forecast month t^* , we construct features \mathbf{C}_{t^*} using covariates and past patrol effort. Sampling from the base distribution $\boldsymbol{\psi}_{t^*}^{(0)}$ and integrating the velocity field yields latent logits $\boldsymbol{\psi}_{t^*}^{(1)}$, converted to risk via $\sigma(\cdot)$. Given

planned patrol efforts $\{a_{i,t^*,j}\}$, detection probabilities follow $p_{i,t^*,j} = \sigma(g_\phi(\mathbf{x}_{i,t^*}, a_{i,t^*,j}))$. Note that the encoder is only used during training and not required at inference.

5 Experiments

5.1 Experimental Setup

Data processing. We use ranger patrol and poaching records from Murchison Falls National Park (MFNP, 2014–2021) and Queen Elizabeth National Park (QENP, 2011–2016), provided by the Uganda Wildlife Authority. Following Kong et al. (2025), we focus on historically high-risk subregions by constructing groups of cells each month: we first select the 20 cells with the highest observed poaching counts, then iteratively expand each group by adding adjacent cells until reaching a maximum of 25 cells. To address non-stationarity, we adopt the windowing strategy of Xu et al. (2020), where for a given test year, only the previous three years of data are used for training.

Baselines. For a fair comparison, we use the same linear detection head across all methods, varying only the model used to predict the latent occupancy state. Each baseline is trained by maximizing the log-likelihood of visit-level observations. Specifically, we compare the following models:

- **Logistic Regression (LogReg).** A linear occupancy model optimized jointly with the detection head using the BFGS optimizer (Nocedal and Wright 2006).
- **Gaussian Process (GP).** Trained with the detection head using variational inference with inducing points. The objective combines the expected log-likelihood of observed snares with a KL divergence regularizer toward the prior.
- **Multilayer Perceptron (MLP).** A three-layer neural network trained with the AdamW optimizer (Loshchilov and Hutter 2019).
- **Graph Neural Network (GNN).** Encodes spatial relationships among cells via graph convolutional layers (Kipf and Welling 2017).
- **Transformer.** Uses self-attention layers to propagate information across cells, capturing long-range dependencies and assigning adaptive weights to neighbors, unlike GNNs which aggregate uniformly (Loshchilov and Hutter 2019).
- **Diffusion Model.** Trained in the same two-stage framework as our method, but replacing composite flow matching with a conditional denoising diffusion model.

We provide further experimental details in Appendix A.2.

Metric. Since the latent occupancy state is unobserved, we evaluate performance at the observation level. This aligns with the practical goal: given a fixed patrol effort, the key question is whether rangers can detect any poaching activity in a cell during a given month. We therefore compute the probability of at least one detection as $\hat{p}_{\text{any},i,t} = r_{i,t} \left(1 - \prod_{j=1}^{J_{i,t}} (1 - p_{i,t,j})\right)$, and compare it against the binary monthly label $y_{i,t}$, indicating whether any poaching was observed in that month. We then evaluate predictions using the area under the precision–recall curve (AUPR), which

is particularly appropriate for this task given the severe class imbalance.

5.2 Main Results

Table 1 reports results on MFNP and QENP. We summarize our findings as follows.

First, our method significantly outperforms baselines on both datasets. On MFNP, we achieve clear improvements in 3 out of 5 years and match the strongest baseline in 1 year. On QENP, we achieve 2 clear wins out of 3 years. On average, WILDFLOW improves AUPR by 7.0% on MFNP and 10.2% on QENP compared to the strongest baseline. Second, surprisingly, the linear occupancy model (LogReg) often performs competitively, and in some cases even outperforms deep learning methods. We attribute this to the small and noisy nature of the data, where simpler models are less prone to overfitting. This highlights the need for caution and careful design when applying deep learning models in this domain. Third, diffusion models do not perform well in our setting; we identified two primary reasons. (i) Training is less stable than flow matching, likely due to the stochastic sampling process. (ii) When complex environmental features are included, diffusion model performance gets stuck, highlighting the data-hungry nature of diffusion models. In contrast, our method solves this problem with a more data-efficient composite base distribution.

5.3 Ablation Study

We evaluate the contribution of each model component through ablation studies. As shown in Table 2, removing either the composite base or the detection head leads to a consistent drop in AUPR across all years. These results demonstrate that both components are essential, and their combination yields the best overall performance.

5.4 Case Study

In this subsection, we conduct a case study to examine where WILDFLOW works particularly well. For each cell i at time t , we compute the log loss $l_{i,t}^{\text{model}} = -y_{i,t} \log(\hat{p}_{\text{any},i,t}) - (1 - y_{i,t}) \log(1 - \hat{p}_{\text{any},i,t})$. Log loss is a proper scoring rule (Gneiting and Raftery 2007), which evaluates both accuracy and calibration of probabilistic forecasts.

To assess relative performance, we calculate the log-loss difference relative to LogReg: $l_{i,t}^{\text{diff}} = l_{i,t}^{\text{WILDFLOW}} - l_{i,t}^{\text{LogReg}}$. Smaller values of $l_{i,t}^{\text{diff}}$ indicate that our method achieves lower loss and thus more accurate forecasts. We then compare the environmental and patrol-effort features of cells in the 10th percentile of $l_{i,t}^{\text{diff}}$ (i.e., where our approach most strongly outperforms LogReg) against the feature distribution of all cells in the same test year.

In Figure 3, we observe that WILDFLOW consistently achieves stronger performance on cells with higher levels of adjacent patrol effort across all test years. Neighboring patrol effort can induce displacement effects and create intricate spatial dependencies,² and the superior perfor-

²We intentionally do not display spatial maps of poaching predictions in the paper for security reasons.

Park	Year	LogReg	GP	MLP	GNN	Transformer	Diffusion	Ours
MFNP	2017	0.305	0.278 ± 0.001	0.288 ± 0.017	0.336 ± 0.004	0.314 ± 0.016	0.313 ± 0.021	0.374 ± 0.032
	2018	0.344	0.308 ± 0.002	0.325 ± 0.004	0.260 ± 0.000	0.302 ± 0.064	0.279 ± 0.001	0.377 ± 0.011
	2019	0.406	0.359 ± 0.002	0.388 ± 0.007	0.526 ± 0.015	0.475 ± 0.141	0.362 ± 0.003	0.409 ± 0.009
	2020	0.421	0.401 ± 0.003	0.408 ± 0.004	0.366 ± 0.014	0.327 ± 0.002	0.439 ± 0.023	0.473 ± 0.006
	2021	0.360	0.364 ± 0.011	0.335 ± 0.010	0.430 ± 0.023	0.295 ± 0.085	0.342 ± 0.032	0.423 ± 0.004
	Avg	0.367	0.342	0.349	0.384	0.343	0.347	0.411
QENP	2014	0.119	0.107 ± 0.003	0.103 ± 0.003	0.136 ± 0.013	0.100 ± 0.002	0.102 ± 0.003	0.116 ± 0.009
	2015	0.180	0.156 ± 0.002	0.172 ± 0.006	0.111 ± 0.000	0.173 ± 0.016	0.162 ± 0.028	0.201 ± 0.017
	2016	0.258	0.220 ± 0.005	0.227 ± 0.004	0.238 ± 0.009	0.220 ± 0.019	0.201 ± 0.034	0.299 ± 0.027
	Avg	0.186	0.161	0.167	0.162	0.164	0.155	0.205

Table 1: AUPR comparison on MFNP and QENP. All methods use a linear detection head for a fair comparison. Cells in green indicate a clear winner (best mean and the runner-up lies outside the winner’s std), while yellow denotes a practical tie (within one std of the best). LogReg is trained with deterministic optimization and therefore has no variance.

Year	w/o composite base	w/o detection head	Ours
2017	0.341 ± 0.008	0.278 ± 0.011	0.374 ± 0.032
2018	0.368 ± 0.019	0.265 ± 0.006	0.377 ± 0.011
2019	0.382 ± 0.006	0.301 ± 0.006	0.409 ± 0.009
2020	0.429 ± 0.007	0.350 ± 0.005	0.473 ± 0.006
2021	0.338 ± 0.040	0.290 ± 0.009	0.423 ± 0.004

Table 2: Ablation study on MFNP, evaluated using AUPR.




Figure 3: Case study of two predictive features on MFNP. We compare the median values of each feature across all cells with those from the subset of top 10% cells where WildFlow achieves the largest relative improvement over logistic regression.

mance of our method in these regions supports the intuition that flow matching is well-suited to modeling complex, high-dimensional spatial patterns. We also find that WILDFLOW tends to perform better on cells located closer to patrol posts. A possible explanation is that proximity to patrol posts increases both the frequency and reliability of detections, which our method is better able to exploit. Exploring the interpretability of these spatial patterns remains an interesting direction for future work.

6 Path to Deployment

Our objective is to translate the proposed latent composite flow framework into a field-ready tool that guides patrol al-

location under real operational constraints (limited compute, bandwidth, and staff time). We proceed in two stages: (1) *validate utility* via prospective field pilots, and (2) *deploy at scale* with ongoing monitoring.

Prospective field pilots. We will run pilots in MFNP and Cross River National Park, in collaboration with the Uganda Wildlife Authority and the Nigeria National Park Service, respectively. Following Xu et al. (2020), we will focus on regions that had low patrol effort in the past to reduce bias from prior ranger knowledge. The model gives each cell a risk score, which we place into three groups: high, medium, and low. Rangers receive patrol plans without seeing these risk labels, and we keep patrol effort similar across the three groups. We will measure effectiveness by (i) snares found per km patrolled and (ii) whether higher-risk groups yield more detections than lower-risk groups.

Responsible rollout. Contingent on conclusive field-trial results, we will integrate the tool with SMART conservation software (SMART 2013) for broader use, conduct ranger workshops to incorporate practitioner feedback, and compare model guidance with ranger intuition. Deployment will be staged (pilot → phased scale-up) with a continuous feedback loop to audit and mitigate bias, and to update models on a regular cadence.

7 Conclusion

In this paper, we present WILDFLOW, the first generative AI-based approach for poaching prediction that explicitly addresses practical challenges such as observation bias and data scarcity. WILDFLOW achieves substantial improvements over various baselines on datasets from two national parks in Uganda. These results establish a strong foundation for applying generative AI techniques in wildlife conservation. As next steps, we aim to further reduce inference costs and conduct rigorous field evaluations to assess the framework’s real-world effectiveness, with the ultimate goal of enabling large-scale deployment.

Acknowledgements

We are thankful to the Uganda Wildlife Authority for granting us access to incident data from Murchison Falls and Queen Elizabeth National Park. We also thank the anonymous reviewers for their valuable feedback. This work was supported by ONR MURI N00014-24-1-2742.

References

- Abernethy, K.; Coad, L.; Taylor, G.; Lee, M.; and Maisels, F. 2013. Extent and ecological consequences of hunting in Central African rainforests in the twenty-first century. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1625): 20120303.
- Barbosa, J. S.; Gondhali, U.; Petrossian, G.; Sharma, K.; Chakraborty, S.; Jacquet, J.; and Freire, J. 2025. A Cost-Effective LLM-based Approach to Identify Wildlife Trafficking in Online Marketplaces. *Proceedings of the ACM on Management of Data*, 3(3): 1–23.
- Coad, L.; Fa, J. E.; Abernethy, K.; Van Vliet, N.; Santamaria, C.; Wilkie, D.; El Bizri, H. R.; Ingram, D. J.; Cawthorn, D.-M.; and Nasi, R. 2019. *Towards a sustainable, participatory and inclusive wild meat sector*. CIFOR.
- Critchlow, R.; Plumptre, A. J.; Driciru, M.; Rwetsiba, A.; Stokes, E. J.; Tumwesigye, C.; Wanyama, F.; and Beale, C. 2015. Spatiotemporal trends of illegal activities from ranger-collected data in a Ugandan national park. *Conservation biology*, 29(5): 1458–1470.
- Dancer, A.; Keane, A.; Beale, C. M.; Dobson, A. D.; Amin, R.; Freeman, R.; Imong, I.; Jones, K.; Linkie, M.; Long, B.; et al. 2022. Evidence of deterrence from patrol data: Trialling application of a differenced-CPUE metric. *Conservation Science and Practice*, 4(8): e12746.
- Dobson, A. D.; Milner-Gulland, E.; Aebischer, N. J.; Beale, C. M.; Brozovic, R.; Coals, P.; Critchlow, R.; Dancer, A.; Greve, M.; Hinsley, A.; et al. 2020. Making messy data work for conservation. *One Earth*, 2(5): 455–465.
- Emogor, C. A.; Ingram, D. J.; Coad, L.; Worthington, T. A.; Dunn, A.; Imong, I.; and Balmford, A. 2021. The scale of Nigeria’s involvement in the trans-national illegal pangolin trade: Temporal and spatial patterns and the effectiveness of wildlife trade regulations. *Biological Conservation*, 264: 109365.
- Gao, R.; Zheng, S.; He, J.; and Shen, L. 2020. CycleGAN-Based Image Translation for Near-Infrared Camera-Trap Image Recognition. In *International Conference on Pattern Recognition and Artificial Intelligence*, 453–464. Springer.
- Gholami, S.; Ford, B.; Fang, F.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; and Mabonga, J. 2017. Taking it for a test drive: a hybrid spatio-temporal model for wildlife poaching prediction evaluated through a controlled field test. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 292–304. Springer.
- Gholami, S.; Mc Carthy, S.; Dilkina, B.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; Mabonga, J.; et al. 2018. Adversary models account for imperfect crime data: Forecasting and planning against real-world poachers. International Conference on Autonomous Agents and Multiagent Systems.
- Gibbons, A.; King, E.; Donohue, I.; and Parnell, A. 2024. Generative AI-based data augmentation for improved bioacoustic classification in noisy environments. *arXiv preprint arXiv:2412.01530*.
- Gneiting, T.; and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477): 359–378.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Kar, D.; Ford, B.; Gholami, S.; Fang, F.; Plumptre, A.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; et al. 2017. Cloudy with a chance of poaching: Adversary behavior modeling and forecasting with real-world poaching data.
- Keane, A.; Jones, J. P.; and Milner-Gulland, E. 2011. Encounter data in resource management and ecology: pitfalls and possibilities. *Journal of Applied Ecology*, 48(5): 1164–1173.
- Kim, E.; Moon, J.; Shim, J.; and Hwang, E. 2023. DualDiscWaveGAN-based data augmentation scheme for animal sound classification. *Sensors*, 23(4): 2024.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.
- Kong, L.; Wang, H.; Pan, Y.; Kim, C. W.; Song, M.; Nguyen, A.; Wang, T.; Xu, H.; and Tambe, M. 2025. Robust optimization with diffusion models for green security. *arXiv preprint arXiv:2503.05730*.
- Lipman, Y.; Chen, R. T. Q.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. *International Conference on Learning Representations (ICLR)*.
- MacKenzie, D. I.; Nichols, J. D.; Lachman, G. B.; Droege, S.; Andrew Royle, J.; and Langtimm, C. A. 2002. Estimating site occupancy rates when detection probabilities are less than one. *Ecology*, 83(8): 2248–2255.
- Manrique, J. S. Ñ.; Gómez, F.; and Hernández-Romero, F. 2025. Anuran call synthesis with diffusion models for enhanced bioacoustic classification under data scarcity. *Ecological Informatics*, 103322.
- Nguyen, T. H.; Sinha, A.; Gholami, S.; Plumptre, A.; Joppa, L.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; and Critchlow, R. 2016. Capture: A new predictive anti-poaching tool for wildlife protection.
- Nocedal, J.; and Wright, S. J. 2006. *Numerical optimization*. Springer.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Rytwinski, T.; Muir, M. J.; Miller, J. R.; Smith, A.; Kelly, L. A.; Bennett, J. R.; Öckerman, S. L.; Taylor, J. J.; Lemieux, A. M.; Pickles, R. S.; et al. 2024. What is the evidence that counter-wildlife crime interventions are effective for conserving African, Asian and Latin American wildlife directly threatened by exploitation? A systematic map. *Ecological Solutions and Evidence*, 5(2): e12323.

SMART. 2013. Spatial Monitoring and Reporting Tool (SMART). <http://smartconservationtools.org/>. Accessed: 2025-08-19.

Sohl, H. 2022. Troubling Trends in Tiger Trade: Q&A with Heather Sohl. https://tigers.panda.org/news_and_stories/stories/troubling_trends_in_tiger_trade_qa_with_heather_sohl/. Accessed: 2025-11-16.

Thiault, L.; Weekers, D.; Curnock, M.; Marshall, N.; Pert, P. L.; Beeden, R.; Dyer, M.; and Claudet, J. 2020. Predicting poaching risk in marine protected areas for improved patrol efficiency. *Journal of Environmental Management*, 254: 109808.

United Nations Office on Drugs and Crime. 2020. World Wildlife Crime Report 2020: Pangolins. Technical report, United Nations Office on Drugs and Crime, Vienna.

United Nations Office on Drugs and Crime. 2024. World Wildlife Crime Report 2024: Trafficking in Protected Species. Technical report, United Nations Office on Drugs and Crime, Vienna.

Wittemyer, G.; Northrup, J. M.; Blanc, J.; Douglas-Hamilton, I.; Omondi, P.; and Burnham, K. P. 2014. Illegal killing for ivory drives global decline in African elephants. *Proceedings of the National Academy of Sciences*, 111(36): 13117–13121.

Xu, L.; Gholami, S.; McCarthy, S.; Dilkina, B.; Plumptre, A.; Tambe, M.; Singh, R.; Nsubuga, M.; Mabonga, J.; Driciru, M.; et al. 2020. Stay ahead of Poachers: Illegal wildlife poaching prediction and patrol planning under uncertainty with field test evaluations (Short Version). In *2020 IEEE 36th international conference on data engineering (ICDE)*, 1898–1901. IEEE.

Xu, L.; Perrault, A.; Fang, F.; Chen, H.; and Tambe, M. 2021. Robust reinforcement learning under minimax regret for green security. In *Uncertainty in Artificial Intelligence*, 257–267. PMLR.

Yang, R.; Ford, B.; Tambe, M.; and Lemieux, A. 2014. Adaptive resource allocation for wildlife protection against illegal poachers. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 453–460.

Zhang, Q.; Yi, X.; Guo, J.; Tang, Y.; Feng, T.; and Liu, R. 2023. A few-shot rare wildlife image classification method based on style migration data augmentation. *Ecological Informatics*, 77: 102237.