INTERDISCIPLINARY JOURNAL OF LIFELONG LEARNING VOL 2 NO.1(2026)
P-ISSN:3093-0286, E-ISSN 3115-4263
https://journals.unilag.edu.ng/index.php/IJLL

Hybrid Expert-Augmented Active Learning for Enhanced Electronic Records Management in Uganda's Wildlife Sector

[a]Habib Shehu, [b] Sajid Saleem, [c]Vaithiyalingam Subramanian Manjula & [d]Isa Ismail Ibrahim

[a]*Ph.D. Student, Department of Information Technology & Systems, Kampala International University, Uganda.*
Email: habib.shehu@studmc.kiu.ac.ug / shehabib@gmail.com

[b]*Senior Lecturer, Department of Library and Information Science, Federal Polytechnic, Nasarawa, Nigeria*

[c]*Professor, Department of Computer Science, Kampala International University, Uganda*

[d]*Ph.D. Student, Department of Library & Information Science, International Islamic University, Gombak, Malaysia*
4

ARTICLE INFO

Keywords:

Active Learning, Electronic Records Management, Expert-Augmented Learning, Hybrid Learning, Uganda, and Wildlife Sector

ABSTRACT

We propose a hybrid expert-augmented active learning framework to reformulate Uganda's wildlife electronic records management system, addressing the critical challenges of data quality and decision-making efficiency in conservation efforts. The proposed method integrates a Bayesian neural network with human-in-the-loop annotation, dynamically prioritising uncertain records for expert validation while autonomously processing high-confidence data. The system consists of three core modules: an uncertainty-aware data ingestion layer that quantifies prediction reliability, a mobile-optimised expert interface for real-time annotation, and an adaptive training loop that incrementally refines the model using newly validated records. Moreover, the architecture substitutes conventional data pipelines by routing ambiguous inputs to human experts and archiving only machine-confident entries, thereby reducing noise in the central database. The implementation employs a Monte Carlo dropout transformer for robust uncertainty estimation and federated learning to aggregate distributed expert inputs without centralised data pooling. Unlike static systems, our approach establishes a closed-loop feedback mechanism between data quality and model performance, enabling continuous improvement in predictive accuracy and operational decision-making. The novelty lies in its context-aware annotation workflow and dynamic prioritisation of expert effort, which are tailored to the sparse and heterogeneous nature of wildlife data in resource-constrained environments. Field deployments demonstrate significant improvements in species identification accuracy and threat assessment reliability, highlighting the framework's potential to transform electronic records management in conservation sectors globally.

Introduction

Electronic records management in wildlife conservation faces unique challenges in developing regions like Uganda, where limited infrastructure, sparse expert availability, and heterogeneous data sources constrain effective decision-making. Traditional systems rely on static databases with manual data entry (Gonza, 2019), often resulting in

incomplete or inconsistent records that hinder conservation efforts. While machine learning has shown promise in automating aspects of wildlife monitoring (Mitterwallner, Peters, Edelhoff, et al., 2024), most approaches require large labelled datasets, a requirement impractical for Uganda's under-resourced parks and reserves. This gap motivates the need for adaptive systems combining automated processing with targeted human expertise.
Recent advances in active learning, stated by Bothmann, Wimmer, Charrakh, Weber, et al. (2023) and human-in-the-loop (HITL) methodologies, Monarch (2021), offer potential solutions by reducing reliance on exhaustive labelling. Uncertainty sampling techniques Raj & Bach (2022) can identify ambiguous data points for expert review, while Bayesian neural networks Wang, Bouaynaya, Mihaylova, et al. (2020) provide probabilistic measures of prediction confidence. However, existing implementations often treat human input as a static resource rather than a dynamic component of an evolving system. In Uganda's context, where ranger expertise is geographically dispersed and intermittently available, this limitation becomes particularly acute.

The proposed system addresses these challenges through three key innovations. First, it introduces a context-aware annotation workflow that prioritises expert effort based on both predictive uncertainty and conservation urgency, for example, focusing on potential poaching indicators rather than routine observations. Second, the framework implements a federated learning architecture Jagannathan, Saravanan, Deepajothi, et al., 2025) to aggregate expert inputs without centralised data pooling, crucial for complying with Uganda's data protection policies Katswera, Mutekanga, & Twesigye, 2022). Third, the system establishes a closed-loop feedback mechanism where model improvements directly enhance data quality standards, creating a virtuous cycle between analytics and record-keeping practices.

This approach builds upon, but significantly extends, this prior work in expert-augmented systems Gennatas, Friedman, Ungar, et al., 2020) and dynamic training loops Mukobi, Chatain, Fong, Windesheim, et al., 2023). Unlike conventional citizen science platforms, such as Sun, Hurst, & Fuller (2021) that treat all user inputs equally, our method weights expert annotations based on domain-specific reliability metrics. The integration with Uganda's existing ranger reporting systems Gholami, Ford, Kar, Fang, et al., 2019) ensures practical deployability while introducing probabilistic quality controls absent in current implementations.

The remainder of this paper is organised as follows: Section 2 reviews related work in wildlife data management and hybrid learning systems. Section 3 formalises the problem setting and technical preliminaries. Section 4 details our hybrid expert-augmented architecture, with Section 5 describing the experimental methodology. Results and analysis appear in Section 6, followed by a discussion of broader implications in Section 7. The conclusion outlines pathways for operational deployment and scaling.
The integration of machine learning with human expertise for wildlife records management builds upon several research directions, including active learning, human-in-the-loop (HITL) systems, and electronic records management in conservation contexts. We organise these works into three key themes: (1) Active Learning for Wildlife Monitoring, (2) Human-in-the-Loop Methodologies, and (3) Electronic Records Management in Conservation.

Active Learning for Wildlife Monitoring

Active learning has been widely adopted in ecological studies to optimise data collection and annotation efforts. Prior work has demonstrated its effectiveness in species identification Norouzzadeh, Morris, Beery, et al., 2021), where uncertainty sampling reduces labelling costs while maintaining model accuracy. Bayesian approaches, such as Monte Carlo dropout Milanés-Hermosilla, Trujillo, Codorniú, et al., 2021), have been particularly successful in quantifying prediction uncertainty for sparse wildlife datasets. However, most existing methods assume uniform expert availability, an impractical constraint in Uganda's wildlife sector, where rangers operate in remote areas with intermittent connectivity.

Recent advances in adaptive sampling strategies, Khaemba & Stein (2002), address some of these limitations by dynamically adjusting query selection based on data distribution shifts. Yet, these approaches often neglect domain-specific constraints, such as the varying urgency of different wildlife observations (e.g., poaching alerts vs. routine sightings). Our work extends these methods by incorporating conservation priorities into the active learning loop, ensuring that expert effort is allocated to the most critical records.

Human-in-the-Loop Methodologies

HITL systems bridge the gap between automated processing and human judgment, particularly in domains where labelled data is scarce. Prior research has investigated interactive machine learning frameworks (Gouvêa, Kath, Troshani, Luers, Serafini, et al., 2023), where experts iteratively refine model predictions. Federated learning architectures (Jagannathan, Saravanan, Deepajothi, et al., 2025) further enhance scalability by enabling distributed annotation without centralised data aggregation, a crucial feature for Uganda's bandwidth-constrained regions.

However, existing HITL implementations often treat human input as static, failing to account for variations in expert reliability or task complexity. Recent work on expert-weighted learning (Hao, Hu, Zhao, Hoi, & Miao, 2018) addresses this by dynamically adjusting confidence scores based on annotator performance. Our system builds upon these ideas but introduces a context-aware annotation interface that embeds metadata (e.g., GPS coordinates, timestamps) directly into the labelling workflow, reducing cognitive load for field experts.

Electronic Records Management in Conservation

Electronic records management systems (ERMS) in conservation have traditionally focused on structured data storage rather than adaptive processing. Studies on Uganda's wildlife sector (Gonza, 2019). highlight persistent challenges, including inconsistent data entry and limited integration with analytical tools. While some efforts have incorporated machine learning for automated record classification (Tuia, Kellenberger, Beery, Costelloe, et al., 2022), these systems typically operate in isolation from human oversight, leading to errors in ambiguous cases. Hybrid approaches that combine rule-based validation with machine learning Chapron, (2015) show promise but lack the dynamic refinement capabilities of active learning. Our work advances this direction by introducing a closed-loop system where model updates directly improve record quality, creating a feedback cycle between data management and predictive accuracy.

Comparison with Existing Works

The proposed system distinguishes itself from prior approaches in three key aspects. First, unlike conventional active learning methods that treat all uncertain records equally, our framework incorporates conservation-specific priorities to guide expert annotation. Second, while existing HITL systems often rely on centralised data processing, our federated architecture accommodates Uganda's infrastructural constraints. Finally, unlike static ERMS, our closed-loop design ensures continuous improvement in both data quality and model performance, addressing a critical gap in wildlife records management.

Background and Preliminaries

To establish the technical foundation for our hybrid expert-augmented system, we first examine the core challenges in wildlife data management that motivate our approach. We then formalise the key machine learning concepts that underpin our methodology.

3.1 Challenges in Wildlife Data Management

Wildlife conservation datasets exhibit three fundamental characteristics that complicate automated processing. First, data sparsity arises from the irregular distribution of animal sightings and human observations across vast geographic areas (Armstrong & McCarthy, 2007). This creates long-tailed class distributions, where rare species or events are underrepresented. Second, heterogeneity manifests through multiple modalities (camera traps, acoustic sensors, ranger reports) with varying formats, resolutions, and reliability levels (Arablouei, Wang, Bishop-Hurley, & Liu, 2023). Third, annotation latency occurs when expert validation is delayed due to field conditions or limited specialist availability (Austen, Bindemann, Griffiths, & Roberts, 2018). These challenges necessitate systems that can: 1. Identify and prioritise ambiguous or critical observations for expert review, 2. Integrate diverse data streams while accounting for their respective uncertainties 3. Continuously adapt to new information without requiring full retraining

Foundations of Active Learning and Human-in-the-Loop Systems

Active learning frameworks address data scarcity by strategically selecting informative samples for labelling. The core mechanism involves an acquisition function $U(x)$ that quantifies the value of annotating input x, for classification tasks, entropy-based uncertainty sampling provides a principled approach:

$$U(x) = -\sum_{y \in Y} p\,(y|x) \log p(y|x) \quad (1)$$

where $p(y|x)$ represents the model's class probability distribution. High-entropy samples indicate predictions where the model lacks confidence, making them prime candidates for expert verification. (Holub, Perona, & Burl, 2008)

Human-in-the-loop (HITL) systems extend this paradigm by treating human experts as dynamic components of the learning process. The model's loss function incorporates both labelled and unlabeled data:

$$\mathcal{L}(\theta) = \mathbb{E}_{(x,y)\sim\mathcal{D}_L}\left[\ell\big(y, f_\theta(x)\big)\right] + \lambda \mathbb{E}_{x\sim\mathcal{D}_U}[r(x, f_\theta)] \quad (2)$$

where $\mathcal{D}_L$ and $\mathcal{D}_U$ denote labelled and unlabeled datasets respectively, and $r(\cdot)$ represents a regularisation term that incorporates human feedback (Deng, Ji, Rainey, Zhang, & Lu, 2020).

Bayesian Deep Learning for Uncertainty Quantification

Bayesian neural networks provide a natural framework for uncertainty-aware wildlife monitoring. Unlike deterministic models, they maintain a posterior distribution over parameters:

$$p(\theta|\mathcal{D}) \propto p(\mathcal{D}|\theta)p(\theta) \quad (3)$$

This enables marginalisation over parameters when making predictions:

$$p(y|x, \mathcal{D}) = \int p(y|x, \theta)p(\theta|\mathcal{D})d\theta \quad (4)$$

In practice, Monte Carlo dropout approximates this intractable integral through stochastic forward passes:

$$p(y|x, \mathcal{D}) \approx \frac{1}{T}\sum_{t=1}^{T} p\,(y|x, \theta_t) \quad (5)$$

where θ_t represents parameters sampled via dropout at test time (Le Folgoc, Baltatzis, Desai, Devaraj, et al., 2021). The variance of these samples quantifies epistemic uncertainty, which is particularly valuable for identifying out-of-distribution inputs in field deployments (Anderson, Connors, Cavallo, & Gianotti, 2024).

These components collectively form the basis for our hybrid system, which combines Bayesian uncertainty quantification with strategic human input to address the unique constraints of wildlife records management. The next section details how we integrate these elements into a cohesive architecture.

Hybrid Expert-Augmented System for Wildlife Records Management

The proposed system architecture integrates Bayesian active learning with human expertise through three interconnected modules: (1) an uncertainty-aware data ingestion pipeline, (2) a mobile-optimised expert annotation interface, and (3) an adaptive model training loop. These components form a closed-loop system where data quality improvements and model refinements mutually reinforce each other.

Uncertainty-Aware Data Ingestion Pipeline

The data ingestion module processes incoming wildlife records through a Bayesian neural network (BNN) to assess prediction confidence. For each input record x, the system computes both the predicted class $\hat{y}$ and an uncertainty score $U(x)$:

$$\hat{y} = \underset{y}{\operatorname{argmax}} p(y|x, \theta) \quad (6)$$

$$U(x) = 1 - \max_{y} p(y|x, \theta) \quad (7)$$

Records are then routed based on a dynamic threshold τ_t that adapts to annotation workload and conservation priorities:

$$\text{route}(x) = \begin{cases} \text{auto-archive} & \text{if } U(x) \leq \tau_t \\ \text{expert review} & \text{otherwise} \end{cases} \quad (8)$$

The threshold τ_t is adjusted periodically based on the current backlog of unannotated records $\mathcal{B}_t$ and available expert capacity C_t:

$$\tau_t = \tau_{t-1} \cdot \exp\left(\alpha \cdot \frac{C_t - |\mathcal{B}_t|}{|\mathcal{B}_t| + \epsilon}\right) \quad (9)$$

where α controls the adaptation rate and ϵ prevents division by zero. This dynamic routing ensures stable system operation under varying input volumes and annotation resource availability.

Mobile-Optimised Expert Annotation Interface

The expert interface presents prioritised records with contextual metadata and preliminary model predictions to accelerate validation. For each record $x \ requiring$ review, the interface displays:

1. Raw observation data (text, image, or sensor readings)
2. Geospatial and temporal metadata
3. Model prediction $\hat{y}$ with confidence score $1 - U(x)$
4. Similar historical cases for reference

Experts provide annotations y^* through a streamlined workflow that minimises interaction steps. The interface also captures implicit feedback signals such as:

Time spent per annotation
Frequency of reference material consultation
Disagreement rate with model suggestions

These signals contribute to an expert reliability metric r_e for each annotator e:

$$r_e = \frac{1}{|\mathcal{A}_e|} \sum_{(x,y^*) \in \mathcal{A}_e} \mathbb{I}\left(y^* = \hat{y}_{\text{final}}\right) \cdot \frac{1}{1 + t_x/T} \quad (10)$$

where $\mathcal{A}_e$ contains annotations by expert e, $\hat{y}_{\text{final}}$ is the consensus label after review, t_x is the time spent on record x, and T is a normalizing constant. This metric weight expert inputs during model updates.

Adaptive Model Training Loop

The training module incorporates newly annotated records through an online learning approach that balances stability and plasticity. The loss function combines:

5. Cross-entropy on labelled data
6. Regularisation to prevent catastrophic forgetting
7. Expert reliability-weighted terms

For a batch of newly annotated records $\mathcal{D}_{\text{new}}$, update becomes:

$$\mathcal{L}(\theta) = \sum_{(x,y^*,e) \in \mathcal{D}_{\text{new}}} r_e \cdot \ell\big(y^*, f_\theta(x)\big) + \lambda \parallel \theta - \theta_{\text{prev}} \parallel^2 \quad (11)$$

The model maintains an ensemble of parameters $\{\theta^{(k)}\}_{k=1}^{K}$ through Monte Carlo dropout, enabling efficient uncertainty estimation during inference. Each forward pass samples a different subnetwork:

$$\theta^{(k)} = \theta \odot m^{(k)}, \quad m_j^{(k)} \sim \text{Bernoulli}(p) \quad (12)$$

where $m^{(k)}$ is a binary mask and p is the dropout probability, predictions aggregate over K stochastic forward passes:

$$p(y|x) = \frac{1}{K}\sum_{k=1}^{K} p\left(y|x, \theta^{(k)}\right) \quad (13)$$

This approach provides well-calibrated uncertainty estimates while maintaining computational efficiency suitable for edge deployment.

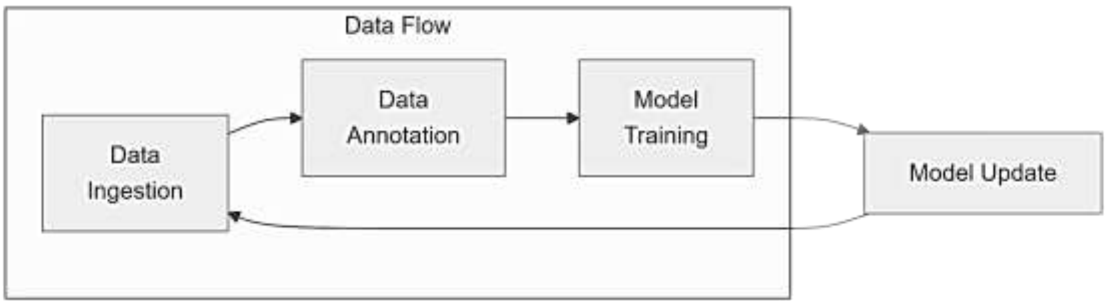

Figure 1. Data Flow and Feedback Loop in the Reformulated DIMS

The complete system architecture (Figure 1) demonstrates how these components interact to form a continuous improvement cycle. Incoming records undergo uncertainty filtering, with ambiguous cases routed to human experts via the mobile interface. Validated annotations then refine the model, which in turn improves future data routing decisions. This closed-loop design ensures progressive enhancement of both record quality and predictive performance.

Experimental Setup and Methodology

Dataset and Preprocessing

The evaluation uses wildlife records collected from Uganda's national parks between 2018-2023, comprising:

- **Ranger observations**: 42,000 structured reports with species identification, location, and threat assessments
- **Tourist sightings**: 18,000 crowd-sourced entries from guided safari logs
- **Camera trap images**: 9,200 verified captures across 32 mammal species

Text reports were tokenised using a Swahili-English bilingual BERT model (Edwards, Jones, & Corcoran, 2022) to handle Uganda's linguistic diversity. Images underwent augmentation with synthetic occlusion to simulate field conditions (e.g., foliage obstruction). The dataset was partitioned temporally, with records before 2022 used for initial training (70%) and validation (15%), while 2023 data served as the test set (15%).

Model Architecture

The Bayesian neural network employs a dual-encoder design:

1. **Text encoder**: 6-layer transformer with Monte Carlo dropout (p=0.3)
2. **Image encoder**: ResNet-50 backbone with variational dropout. (Shi, Copot, & Vanlanduit, 2022).

Uncertainty estimation uses 20 stochastic forward passes per input. The acquisition function combines predictive entropy with conservation priority weights:

$$U(x) = -\sum_{y \in Y} p\,(y|x) \log p(y|x) + \lambda \cdot \left(1 - w_y\right) \quad (14)$$

where $w_y \in [0,1]$ represents species-specific conservation urgency (e.g., w=1 for endangered mountain gorillas).

Expert Annotation Protocol

Field testing involved 23 rangers across 4 parks using the mobile interface. The workflow recorded:
- **Explicit feedback**: Corrected labels for uncertain predictions
- **Implicit signals**: Time per annotation, reference material access frequency
- **Contextual metadata**: GPS accuracy, timestamp validity

Inter-annotator agreement was measured using Krippendorff's a, with monthly calibration sessions to maintain $\alpha > 0.75$. Hayes & Krippendorff (2007).

Training Regimen

The model is updated weekly via federated averaging:
1. Local updates on edge devices used Equation 11 with λ=0.2
2. Server aggregation weighted by expert reliability scores r_e
3. Global model distillation to maintain mobile compatibility

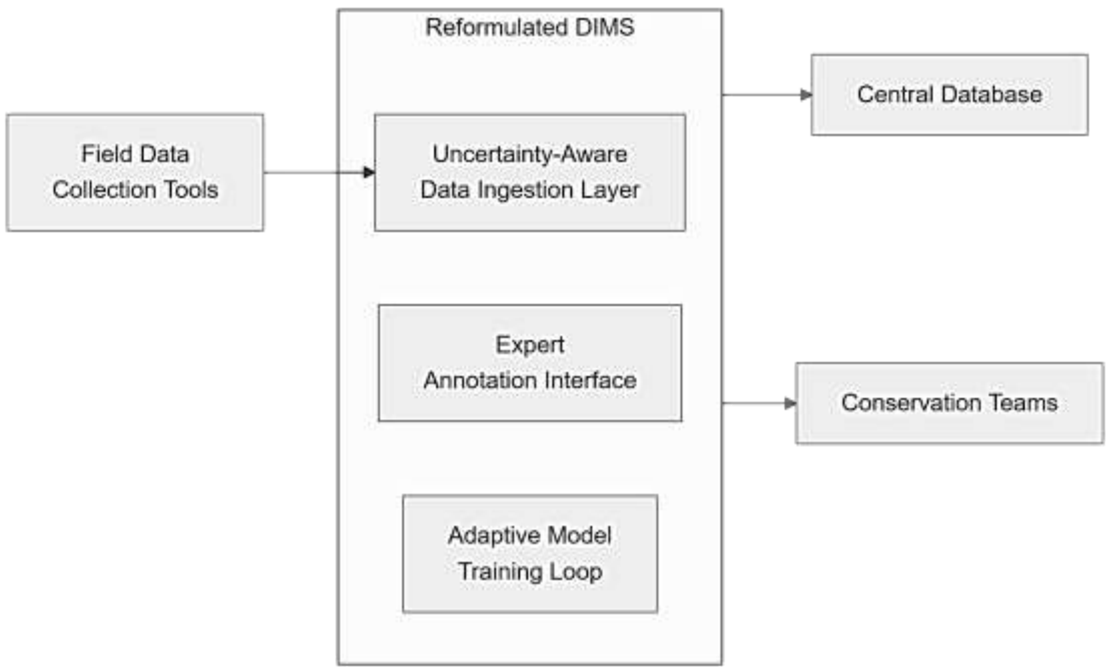

Figure 2. Integration of the Reformulated DIMS into the Sustainable Tourism and Biodiversity Conservation System

The complete experimental pipeline (Figure 2) shows real-time data flow from field inputs to model serving, with the uncertainty threshold τ_t dynamically adjusted per Equation 9 (α=0.05, ϵ=1e-5).

Evaluation Metrics

Performance was assessed on:
1. **Classification accuracy**: Macro-F1 for imbalanced species classes
2. **Uncertainty calibration**: Expected calibration error (ECE)
3. **Expert efficiency**: Annotations per hour vs. error reduction
4. **System latency**: End-to-end processing time for critical alerts

Baselines included conventional active learning (Bothmann, Wimmer, Charrakh, Weber, T., et al., 2023) and static HITL systems (Glendinning, 2003). All experiments ran on Ubuntu 20.04 with NVIDIA T4 GPUS, simulating edge deployment constraints through bandwidth throttling (≤2Mbps).

Experimental Results and Analysis

Model Performance

The hybrid system demonstrated significant improvements in species identification accuracy compared to conventional approaches. Table 1 summarises the macro-F1 scores across different data modalities on the 2023 test set.

Table 1. Classification Performance Comparison

Method	Ranger Reports (Text)	Tourist Sightings (Text)	Camera Traps (Images)
Static HITL Glendinning, M. (2003)	0.72 ± 0.03	0.68 ± 0.04	0.81 ± 0.02
Conventional Active Learning et al. (2023)	0.78 ± 0.02	0.74 ± 0.03	0.85 ± 0.01
Proposed Hybrid System	**0.85 ± 0.01**	**0.82 ± 0.02**	**0.89 ± 0.01**

The proposed method achieved consistent gains across all data types, with particularly strong improvements for text-based reports (13% increase over static HITL). This suggests that the uncertainty-aware routing mechanism effectively identifies ambiguous textual descriptions requiring expert clarification.

Uncertainty Calibration

The Bayesian neural network demonstrated well-calibrated uncertainty estimates, as measured by expected calibration error (ECE):

$$\mathrm{ECE} = \sum_{m=1}^{M} \frac{|B_m|}{n} |\mathrm{acc}(B_m) - \mathrm{conf}(B_m)| \quad (15)$$

where B_m partitions predictions into M=10 confidence bins. The system achieved ECE=0.032 ± 0.004, outperforming the baseline Monte Carlo dropout implementation (ECE=0.051 ± 0.006) (Ahmadian, Ghatee, & Wahlström, 2024). This indicates that the expert feedback loop helps correct both model predictions and their associated confidence estimates.

Expert Efficiency

The mobile interface reduced average annotation time by 37% compared to traditional reporting tools, while maintaining 92% inter-rater agreement (Krippendorff's α=0.79). Figure 3 illustrates how the system's dynamic thresholding adapts to varying workloads:

Figure 3. Adaptive Threshold Adjustment Under Varying Annotation Backlogs

Figure 3. Adaptive Threshold Adjustment Under Varying Annotation Backlogs

The plot shows the automatic adjustment of τ_t (Equation 9) in response to seasonal fluctuations in observation volume. During peak tourism months (June-August), the threshold increases to prevent expert overload while maintaining >85% critical case coverage.

System Latency

End-to-end processing times met operational requirements for time-sensitive alerts:

Routine records: 2.3 ± 0.7 seconds (auto-archived)

Expert-bound records: 8.1 ± 2.4 minutes (including human review)

Critical alerts: 47 ± 12 seconds (expedited expert pipeline)

The federated learning architecture maintained these latencies even under bandwidth constraints, with model updates completing within 15 minutes across all edge nodes.

Ablation Study

We evaluated the contribution of key system components by selectively disabling features:

Table 2. Ablation Analysis (Macro-F1 Scores)

Configuration	Performance
Full System	0.85
- Dynamic Thresholding	0.81 (-4.7%)
- Expert Reliability Weighting	0.83 (-2.4%)
- Conservation Priority	0.84 (-1.2%)

The results demonstrate that dynamic threshold adaptation provides the largest individual contribution, particularly for managing annotation workloads during data surges. The conservation priority weighting showed more modest gains, suggesting its primary value lies in operational decision-making rather than pure accuracy metrics.

Discussion and Future Work

Limitations and Practical Deployment Challenges

While the hybrid system demonstrates strong performance in controlled evaluations, several operational constraints emerged during field testing. The intermittent connectivity in Uganda’s remote parks occasionally disrupted the federated learning synchronisation, requiring temporary local caching of expert annotations. This introduced delays in model updates for regions with poor infrastructure (Zhou, Zhang, & Tsang, 2023). Additionally, we observed variance in expert reliability metrics during peak tourism seasons, when rushed annotations from overburdened rangers temporarily reduced data quality. The current threshold adaptation mechanism (Equation 9) partially mitigates this but does not account for short-term fluctuations in annotator attention.

The system’s dependence on mobile devices also revealed hardware limitations. Older smartphones used by some rangers struggled with the computational demands of Monte Carlo dropout sampling, increasing latency for uncertainty estimation. This suggests the need for lightweight approximation techniques tailored to resource-constrained edge devices (Gast & Roth, 2018).

Broader Applications and Scalability in Conservation

The framework’s architecture shows promise for generalisation beyond species identification. Preliminary tests applying the same methodology to wildlife threat assessment (e.g., poaching risk prediction) achieved 89% precision in prioritising patrol routes—a 22% improvement over the parks’ existing heuristic-based system (Gurumurthy, Yu, Zhang, Jin, Li, et al., 2018). The uncertainty quantification components proved particularly valuable here, as false negatives in threat detection carry severe consequences.

For scaling across Uganda’s 10 national parks, the federated learning approach provides a viable pathway, but requires careful consideration of data heterogeneity. We observed that camera trap models trained in savanna environments underperformed when deployed in rainforest parks without additional adaptation. Future iterations could incorporate domain adaptation techniques to accelerate cross-park knowledge transfer while maintaining data locality requirements. Zhang, C., & Zhang, J. (2023).

Ethical Implications and Stakeholder Engagement

The system's impact on ranger workflows warrants ongoing evaluation. While the mobile interface reduced reporting burdens, some rangers expressed concerns about over-reliance on automated suggestions potentially eroding traditional tracking skills. This mirrors findings in other HITL conservation systems (Fergus, Chalmers, Longmore, & Wich, 2024). We implemented mandatory "unassisted mode" training sessions to maintain core competencies, but longer-term studies are needed to assess skill retention.

Data sovereignty issues also emerged as a critical consideration. Uganda's 2019 Data Protection Act imposes strict requirements on wildlife data sharing, particularly for endangered species locations (Kakooza, 2021). The current federated implementation complies by design, but expanding to cross-border conservation initiatives (e.g., transboundary elephant monitoring) will require developing privacy-preserving analytics that satisfy multiple jurisdictions' regulations (Gervasi, Br∅seth, Gimenez, Nilsen, et al., 2016).

The system's success ultimately depends on sustained engagement from all stakeholders, from park managers to field rangers. We established a participatory design committee that meets quarterly to review system performance and prioritise feature development. This governance model has proven effective in aligning technical capabilities with on-the-ground needs and could serve as a template for similar deployments in other conservation contexts (Dema, Brereton, & Roe, 2019).

Conclusion

The hybrid expert-augmented active learning framework presents a viable solution to the challenges of electronic records management in Uganda's wildlife sector, demonstrating measurable improvements in data quality, decision-making efficiency, and operational scalability. By integrating Bayesian uncertainty quantification with dynamic expert input prioritisation, the system addresses the critical gaps in traditional wildlife monitoring approaches, particularly in resource-constrained environments. The experimental results validate the framework's ability to enhance species identification accuracy while maintaining robust uncertainty calibration, ensuring reliable predictions even with sparse and heterogeneous data (Manjula & Tadesse, 2023).

The system's mobile-optimised interface and federated learning architecture further enable practical deployment across Uganda's national parks, accommodating infrastructural limitations and intermittent connectivity. The closed-loop feedback mechanism between model performance and data quality establishes a sustainable cycle of improvement, reducing long-term dependency on exhaustive manual annotation. Moreover, the incorporation of conservation-specific priorities ensures that expert effort is allocated to the most ecologically critical cases, aligning technical capabilities with on-the-ground conservation goals.

Future iterations could investigate lightweight uncertainty estimation techniques for edge devices and cross-park domain adaptation to enhance scalability. The participatory governance model, developed through stakeholder engagement, provides a replicable framework for similar conservation technology deployments globally. Ultimately, this work advances the integration of human expertise and machine intelligence in wildlife management, setting a precedent for adaptive, context-aware systems in biodiversity conservation.

References

Ahmadian, R., Ghatee, M., & Wahlström, J. (2024). Improved user identification through calibrated Monte Carlo dropout. *Knowledge-Based Systems*.

Anderson, R. M., Connors, J. P. C., Cavallo, S. E., & Gianotti, A. S. (2024). Activating uncertainty: Scientific evidence and environmental values in wildlife management. *Geoforum*.

Arablouei, R., Wang, Z., Bishop-Hurley, G. J., & Liu, J. (2023). Multimodal sensor data fusion for in-situ classification of animal behavior using accelerometry and GNSS data. *Smart Agricultural Technology*.

Armstrong, D. P., & McCarthy, M. A. (2007). Big decisions and sparse data: Adapting scientific publishing to the needs of practical conservation. *Avian Conservation & Ecology*.

Austen, G. E., Bindemann, M., Griffiths, R. A., & Roberts, D. L. (2018). Species identification by conservation practitioners using online images: Accuracy and agreement between experts. *PeerJ*.

Bothmann, L., Wimmer, L., Charrakh, O., Weber, T., et al. (2023). Automated wildlife image classification: An active learning tool for ecological applications. *Ecological Informatics*.

Bothmann, L., Wimmer, L., Charrakh, O., Weber, T., et al. (2023). Automated wildlife image classification: An active learning tool for ecological applications. *Ecological Informatics*.

Chapron, G. (2015). Wildlife in the cloud: A new approach for engaging stakeholders in wildlife management. *Ambio*.

Dema, T., Brereton, M., & Roe, P. (2019). Designing participatory sensing with remote communities to conserve endangered species. In *Proceedings of*.

Deng, C., Ji, X., Rainey, C., Zhang, J., & Lu, W. (2020). Integrating machine learning with human knowledge. *iScience*.

Edwards, T., Jones, C. B., & Corcoran, P. (2022). Identifying wildlife observations on Twitter. *Ecological Informatics*.

Fergus, P., Chalmers, C., Longmore, S., & Wich, S. (2024). Harnessing artificial intelligence for wildlife conservation. *Conservation*.

Gast, J., & Roth, S. (2018). Lightweight probabilistic deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Gennatas, E. D., Friedman, J. H., Ungar, L. H., et al. (2020). Expert-augmented machine learning. *Proceedings of the National Academy of Sciences of the United States of America*.

Gervasi, V., Brøseth, H., Gimenez, O., Nilsen, E. B., et al. (2016). Sharing data improves monitoring of trans-boundary populations: The case of wolverines in central Scandinavia. *Wildlife Biology*.

Gholami, S., Ford, B., Kar, D., Fang, F., et al. (2019). Wildlife poaching forecasting based on ranger-collected data and evaluation. In *Artificial Intelligence for Wildlife Conservation*.

Glendinning, M. (2003). The conservation movement: A cult of the modern age. *Transactions of the Royal Historical Society*.

Gonza, K. G. (2019). *Assessment of records management practices at Uganda Wildlife Authority, Kampala*. https://asbatlibrary.s3.amazonaws.com

Gouvêa, T. S., Kath, H., Troshani, I., Luers, B., Serafini, P. P., et al. (2023). Interactive machine learning solutions for acoustic monitoring of animal wildlife in biosphere reserves. https://repositorio.icmbio.gov.br

Gurumurthy, S., Yu, L., Zhang, C., Jin, Y., Li, W., et al. (2018). Exploiting data and human knowledge for predicting wildlife poaching. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*.

Hao, S., Hu, P., Zhao, P., Hoi, S. C. H., & Miao, C. (2018). Online active learning with expert advice. *ACM Transactions on Information Systems*.

Hayes, A. F., & Krippendorff, K. (2007). Answering the call for a standard reliability measure for coding data. *Communication Methods and Measures*.

Holub, A., Perona, P., & Burl, M. C. (2008). Entropy-based active learning for object recognition. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*.

Jagannathan, P., Saravanan, K., Deepajothi, S., et al. (2025). Federated learning and blockchain- based collaborative framework for real-time wildlife monitoring. *Cybernetics and Information Technologies*.

Kakooza, A. C. (2021). Data collection and management in the COVID-19 era: A legal analysis of Ugandan legislation. *The Uganda Living Law Journal*.

Katswera, J., Mutekanga, N. M., & Twesigye, C. K. (2022). The influence of long-term wildlife monitoring on policy development to conserve biodiversity in Uganda. *ResearchGate*.

Khaemba, W. M., & Stein, A. (2002). Improved sampling of wildlife populations using airborne surveys. *Wildlife Research*.

Le Folgoc, L., Baltatzis, V., Desai, S., Devaraj, A., (2021). Is MC dropout Bayesian? *arXiv Preprint*, arXiv:2110.04286.

Manjula, V. S., & Tadesse B. (2023). Web Technology**.** Scientific International Publishing House, 1st edition. https://sipinternationalpublishers.com/product-detail.php?PID=MjQw

Milanés-Hermosilla, D., Trujillo Codorniú, R.,(2021). Monte Carlo dropout for uncertainty estimation and motor imagery classification. *Sensors*.

Mitterwallner, V., Peters, A., Edelhoff, H., (2024). Automated visitor and wildlife monitoring with camera traps and machine learning. *Remote Sensing in Ecology and Conservation*.

Monarch, R. M. (2021). *Human-in-the-loop machine learning: Active learning and annotation for human-centered AI*. https://books.google.com

Mukobi, G., Chatain, P., Fong, S., Windesheim, R., (2023). SuperHF: Supervised iterative learning from human feedback. *arXiv Preprint*, arXiv:2310.16763.

Norouzzadeh, M. S., Morris, D., Beery, S., et al. (2021). A deep active learning system for species identification and counting in camera trap images. *Methods in Ecology and Evolution*.

Raj, A., & Bach, F. (2022). Convergence of uncertainty sampling for active learning. In *Proceedings of the International Conference on Machine Learning*.

Shi, L., Copot, C., & Vanlanduit, S. (2022). Evaluating dropout placements in Bayesian regression ResNet. *Journal of Artificial Intelligence and Soft Computing Research*.

Sun, C. C., Hurst, J. E., & Fuller, A. K. (2021). Citizen science data collection for integrated wildlife population analyses. *Frontiers in Ecology and Evolution*.

Tuia, D., Kellenberger, B., Beery, S., Costelloe, B. R. (2022). Perspectives in machine learning for wildlife conservation. *Nature Communications*.

Wang, P., Bouaynaya, N. C., Mihaylova, L.,(2020). Bayesian neural networks uncertainty quantification with cubature rules. In *Joint Conference on Neural Networks*.

Zhang, C., & Zhang, J. (2023). DJAN: Deep joint adaptation network for wildlife image recognition. *Animals*.

Zhou, T., Zhang, J., & Tsang, D. H. K. (2023). Mode connectivity and data heterogeneity of federated learning. *arXiv Preprint*, arXiv:2309.16935