Continuous Learning for AI-Informed Landmine Risk Assessment

Nirmal Boghara [1] **Mateo Dulce Rubio** [1]

Abstract

Humanitarian demining efforts aim to assist 60 million people living in landmine-contaminated areas across 70 countries, with complete clearance requiring over 1,100 years at current rates. Recent AI-informed risk assessment systems seek to improve resource allocation but operate statically, providing fixed contamination maps that remain unchanged throughout operations and overlooking the dynamic nature of mine action, where field surveys continuously reveal ground truth that could refine risk assessments. We address this gap by conducting the first systematic evaluation of continuous learning strategies for landmine risk assessment using real operational data from Colombia. We model demining operations as a sequential resource allocation problem under budget constraints and evaluate six learning strategies, comparing passive baselines against retraining approaches, active learning, and active search methods. Our evaluation reveals that straightforward greedy retraining achieves 98.7% average recall under 50% query budget, substantially outperforming competing approaches, while active search methods achieve competitive performance when no historical training data is available. Our results demonstrate the value of model adaptation in humanitarian AI applications under geographic distribution shifts and inform ongoing deployment efforts for continuous learning systems in operational mine action contexts.

1. Introduction

Landmines and other explosive remnants of war contaminate vast territories affected by conflict, representing a persistent lethal impediment to post-conflict recovery and restoration long after hostilities cease. Despite ongoing humanitarian mine action efforts, landmine contamination continues to affect over 60 million people across 70 countries worldwide: at current clearance rates, complete global demining would require approximately 1,100 years (UNMAS, 2023). This issue stems largely from a mismatch between the vast areas requiring surveying and the limited resources for demining operations that must balance between potential false negatives that could leave communities at risk and precautionary overestimation of contaminated areas that inefficiently allocates already scarce resources. This underscores the urgent need for more efficient, evidence-based approaches for planning and prioritization in humanitarian mine action.

Recent advances in geospatial machine learning have enabled AI-informed systems to support humanitarian demining operations by analyzing geographic features, demographic variables, and satellite imagery to estimate explosive ordnance contamination risk (Kischelewski et al., 2024). These approaches have demonstrated potential to predict contamination boundaries during initial non-technical surveys and reduce both false alarms and clearance time (Dulce Rubio et al., 2024). However, existing AI-informed risk assessment systems operate statically, providing a single contamination risk map at the beginning of operations that remains unchanged throughout the demining process. This approach overlooks the dynamic nature of demining activities, where information collected during field surveys and ground operations continuously reveals ground truth about contamination patterns. As demining teams confirm or cancel hazardous areas, this field data remains unused for refining risk assessments, potentially leading to misallocated resources. This limitation is particularly pronounced when operating under geographic distribution shifts, where training data comes from a spatial region (e.g., country, province) different from the deployment target context—a common scenario in humanitarian operations where historical data may be sparse or unavailable in newly affected areas.

We address this gap by conducting the first systematic evaluation of continuous learning strategies for landmine risk assessment using real-world operational data. Specifically, we model demining operations as a sequential AI-informed resource allocation problem under budget constraints and evaluate six learning strategies on data from 15 municipalities in Colombia that have completed mine action operations. We compare passive baselines against simple retrain-

[1]Center for Data Science, New York University, New York, NY USA. Correspondence to: Mateo Dulce Rubio <mateo.d@nyu.edu>.

LatinX in AI Workshop @ ICML 2026, Seoul, South Korea. 2026.

ing approaches, active learning methods that optimize for information gain, and active search methods that prioritize predicted positives assuming no initial training data is available. Our evaluation reveals that a straightforward greedy strategy with continuous retraining achieves 98.7% average recall under 50% query budget, substantially outperforming more sophisticated active learning approaches (74.5% recall) while maintaining computational feasibility.

Our contributions are threefold: (1) we provide the first formalization of humanitarian demining as a sequential decision problem and establish an evaluation framework for continuous learning methods using real operational data under realistic constraints, (2) we conduct a systematic empirical evaluation comparing passive baselines, retraining strategies, active learning, and active search methods, showing that simple greedy retraining outperforms competing alternatives while active search achieves competitive performance when no training data is available, and (3) we analyze why different strategies succeed or fail in this domain and translate findings into deployment guidance for implementing continuous learning capabilities in operational mine action contexts in Colombia and other conflict-affected regions.

2. AI for Humanitarian Demining

Humanitarian demining operations clear conflict-affected territories of explosive ordnance so land can be safely reclaimed by communities. Operations begin with *non-technical surveys* that synthesize official records and community reports to define *hazardous areas* for subsequent technical surveys and clearance (GICHD, 2009). Decision-makers face a fundamental resource allocation challenge: vast areas require surveying, but financial, logistical, and safety constraints limit the number of locations that can be physically inspected. This translates into a tradeoff between potential false negatives that could leave communities at risk and precautionary overestimation of contaminated areas that inefficiently allocates already scarce resources.

2.1. Related Work

Recent advances in geospatial machine learning have proposed AI-informed systems to address the demining resource allocation challenge by predicting contamination risk from geographic features, demographic variables, and conflict-related indicators (Kischelewski et al., 2024). For instance, the RELand system (Dulce Rubio et al., 2024; Collins et al., 2024), deployed in Colombia and Afghanistan, shows that machine learning pipelines can accurately predict contamination boundaries during initial assessments, with potential to reduce false alarms and clearance time by 50%. Similar approaches have been explored in Cambodia and other conflict-affected regions (Cirillo et al., 2024). However, existing AI systems for mine action operate statically: models trained on historical data generate a single fixed risk map that remains unchanged throughout operations. This design misses a critical opportunity: demining operations are inherently sequential, and as teams survey locations, they continuously reveal ground truth labels that could inform subsequent inspection decisions. Each label provides signal about local contamination patterns, yet current systems discard this information rather than using it to refine predictions for unsurveyed areas.

The machine learning literature offers two relevant paradigms for learning from sequential feedback. **Active learning** selects informative instances to label in order to improve overall model performance with minimal queries, typically by prioritizing locations where the model exhibits high uncertainty (Settles, 2009). The underlying objective is to minimize expected generalization error once the labeling budget is exhausted, optimizing for long-term model performance across the entire feature space. However, this objective conflicts with mine action priorities, which focus on maximizing immediate recall (mines found) rather than future predictive capabilities. Moreover, under extreme class imbalance (1-2% prevalence), uncertainty sampling often selects ambiguous negative instances rather than missed positives, wasting limited query budget on locations that do not advance the operational goal of finding contaminated areas.

Active search methods offer a more aligned objective by explicitly optimizing for cumulative discovery of rare positive instances (Garnett et al., 2012). These approaches, applied to domains like drug discovery (Jiang et al., 2018) and tenant eviction (Sarkar et al., 2025), use Bayesian optimization to balance exploration and exploitation when seeking rare targets. However, active search methods typically assume minimal or no training data and rely on computationally simple models (e.g., k-nearest neighbors) for tractable posterior updates (Jiang et al., 2017). This design choice limits their ability to leverage rich historical data or complex feature representations available in mine action contexts. Furthermore, the computational overhead of computing combinatorial probabilistic models over large spatial regions can be prohibitive for field deployment, as evidenced by scalability challenges in our experiments.

Our work bridges these paradigms by formalizing mine action as a continuous learning problem that combines rich historical training data (unlike active search) with recall-focused objectives (unlike active learning). We evaluate both simple retraining strategies and sophisticated selection methods under realistic operational constraints: (1) extreme class imbalance (1-2% prevalence), (2) recall-focused objectives rather than overall prediction error, (3) geographic distribution shift between training and deployment regions,

(4) batch querying driven by operational logistics, and (5) computational constraints for field deployment. Understanding which learning strategies work under these constraints is critical for designing deployable systems that can improve humanitarian outcomes in mine action operations.

3. Problem Formulation

We formalize humanitarian demining as a *sequential batch selection problem under budget constraints*, where the objective is to maximize recall (proportion of landmines found) through adaptive querying strategies. Consider a target region divided into N spatial units (grid cells), where each location i is characterized by observable features $X_i \in \mathcal{X} \subset \mathbb{R}^d$ and an unknown binary label $y_i \in \{0, 1\}$ indicating landmine presence. Humanitarian demining organizations must decide which locations to inspect under a strict budget constraint of $B < N$ queries, where each query reveals the true label y_i for the selected location.

We model this as a sequential decision process over T rounds. At each round t:

1. A predictive model $r_t : \mathcal{X} \to [0, 1]$ estimates contamination risk for all unqueried locations based on features and previously observed labels in a training dataset $\mathcal{D}_t$.

2. A selection strategy chooses a batch $\mathcal{B}_t$ of b_t locations to query, subject to $\sum_{t=1}^{T} b_t = B$.

3. Ground truth labels $\{y_i : i \in \mathcal{B}_t\}$ are revealed.

4. The model is (potentially) retrained to r_{t+1} on the expanded labeled set $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(\mathbf{x}_i, y_i) : i \in \mathcal{B}_t\}$.

The objective is to maximize recall, defined as the proportion of total landmines discovered within the budget B:

$$\text{Recall} = \frac{\sum_{i \in \mathcal{Q}} \mathbb{I}(y_i = 1)}{\sum_{i=1}^{N} \mathbb{I}(y_i = 1)}, \tag{1}$$

where $\mathcal{Q} = \bigcup_{t=1}^{T} \mathcal{B}_t$ denotes all queried locations with $|\mathcal{Q}| \leq B$, and $\mathbb{I}(\cdot)$ is the indicator function.

This formulation enables systematic study of continuous learning approaches that update risk estimates as operations progress. Importantly, it generalizes current practice (static models) as a special case while allowing evaluation of adaptive alternatives. Our framework provides the first rigorous basis for comparing how different learning strategies—from simple retraining to strategic batch sampling—perform under realistic operational constraints including extreme class imbalance, distribution shift, budget constraints, and limited computational resources.

4. Methodology

We evaluate three categories of methods for sequential landmine discovery under budget constraints. **Passive baselines** include random sampling and static risk-based selection that do not adapt to field observations, serving as benchmarks. **Continuous learning methods** retrain predictive models as ground truth is revealed, comparing simple greedy selection and active learning frameworks that balance risk with uncertainty and density. Finally, we evaluate **Active Search without training data**, representing deployment scenarios where no prior contamination data is available and the model must learn entirely from field observations. All methods operate under the sequential decision framework defined above, where at each round t, a batch $\mathcal{B}_t$ of locations is selected from the unlabeled pool $\mathcal{U}_t = \{1, ..., N\} \setminus \mathcal{Q}_{t-1}$, with $\mathcal{Q}_{t-1} = \bigcup_{s=1}^{t-1} \mathcal{B}_s$ denoting previously queried locations.

4.1. Baseline Methods

Random Sampling The random baseline selects locations uniformly at random from the unlabeled pool at each round

$$\mathcal{B}_t \sim \text{Uniform}(\mathcal{U}_t, b_t), \tag{2}$$

without leveraging prior information or feedback from previous inspections. This serves as a lower-bound benchmark.

Static Greedy Selection This baseline uses fixed risk scores $\widehat{r}_i$ (e.g., from a pretrained ML model or expert assessments) that remain unchanged throughout operations. At each round, it selects the b_t highest-risk unlabeled locations:

$$\mathcal{B}_t = \underset{\mathcal{B} \subset \mathcal{U}_t, |\mathcal{B}| = b_t}{\arg\max} \sum_{i \in \mathcal{B}} \widehat{r}_i. \tag{3}$$

This represents a purely exploitative (greedy) strategy without adaptation to observed outcomes.

4.2. Continuous Learning Methods

Greedy Retraining This method retrains a supervised classifier at each round using all labeled data accumulated from previous queries: $\mathcal{D}_t = \mathcal{Q}_{t-1} \cup Q_0$, where Q_0 represents the available historical training data. The model estimates contamination risk for each unlabeled location:

$$\widehat{r}_i^{(t)} = \widehat{\mathbb{P}}(Y = 1 \mid X = x_i), \quad i \in \mathcal{U}_t, \tag{4}$$

and selects the top-b_t locations by predicted risk

$$\mathcal{B}_t = \underset{\mathcal{B} \subset \mathcal{U}_t, |\mathcal{B}| = b_t}{\arg\max} \sum_{i \in \mathcal{B}} \widehat{r}_i^{(t)}. \tag{5}$$

This approach adapts to observed outcomes through model retraining but remains purely exploitative in selection.

Explore-Exploit Strategy This baseline balances exploitation and exploration within each batch by partitioning $\mathcal{B}_t$ into two disjoint sets, according to a fixed $\alpha \in [0, 1]$:

$$\mathcal{B}_t = \mathcal{B}_t^{\text{exploit}} \cup \mathcal{B}_t^{\text{explore}}, \tag{6}$$

where $|\mathcal{B}_t^{\text{exploit}}| = \lfloor \alpha b_t \rfloor$ locations are selected as the top-ranked by $\widehat{r}_i^{(t)}$, and the remaining $|\mathcal{B}_t^{\text{explore}}| = b_t - \lfloor \alpha b_t \rfloor$ are sampled uniformly at random from $\mathcal{U} \setminus \mathcal{B}_t^{\text{exploit}}$. The model is then updated according to equation (5) after each batch, allowing for both model adaptation and exploration.

Active Learning with Hybrid Acquisition This method combines multiple acquisition signals to balance uncertainty reduction, representativeness, and prior risk (Settles, 2009). Model uncertainty is estimated using Shannon entropy $H^{(t)}(x_i)$ over an ensemble of classifiers predicting contamination at location i in round t:

$$H^{(t)}(x_i) = -\sum_{c \in \{0,1\}} \widehat{p}_{ic} \log \widehat{p}_{ic}, \quad i \in \mathcal{U}_t, \tag{7}$$

where $\widehat{p}_{ic}$ denotes the fraction of ensemble members predicting class c. Moreover, to encourage representativeness, an information density term is computed:

$$D^{(t)}(x_i) = \frac{1}{|\mathcal{U}_t|} \sum_{j \in \mathcal{U}_t} \text{sim}(\mathbf{x}_i, \mathbf{x}_j), \quad i \in \mathcal{U}_t, \tag{8}$$

where $\text{sim}(\cdot, \cdot)$ measures feature similarity between locations i and j using cosine similarity.

Vote entropy $H^{(t)}(x_i)$ and information density $D^{(t)}(x_i)$ are then combined into a weighted acquisition score

$$A_{\text{info}}^{(t)}(x_i) = (1 - \beta) H^{(t)}(x_i) + \beta D^{(t)}(x_i), \tag{9}$$

where $\beta \in [0, 1]$ controls the uncertainty-density tradeoff.

Finally, estimated risk scores $\widehat{r}_i^{(t)}$ are incorporated through a budget-adaptive weighting scheme:

$$A^{(t)}(x_i) = (1 - w_t) A_{\text{info}}^{(t)}(x_i) + w_t \widehat{r}_i^{(t)}, \tag{10}$$

where w_t increases as budget is utilized, gradually shifting from exploration to exploitation. When positive discoveries are made and spatial coordinates are available, a proximity bonus encourages querying near previously discovered landmines. Final batch selection uses a greedy diversity-aware strategy to avoid redundant queries.

4.3. Active Search Without Training Data

We evaluate a decision-theoretic active search strategy that assumes no historical training data is available and learns entirely from field observations (Garnett et al., 2012). This method explicitly maximizes the expected cumulative number of landmines discovered over the remaining budget through a non-myopic formulation that directly aligns with the operational goal of maximizing recall.

Model Specification. We follow Jiang et al. (2017; 2018) and model the probability of landmine presence $\mathbb{P}(y_i = 1 \mid \mathbf{x}_i)$ using a nonparametric k-nearest neighbors (k-NN) graph constructed over spatial coordinates. This approach reflects the empirical observation that landmines tend to be spatially clustered. Probabilities are estimated by propagating labels from observed locations to unobserved ones based on inverse Euclidean distance. For an unobserved location $i \in \mathcal{U}_t$, the probability is defined as:

$$\widehat{\mathbb{P}}^{(t)}(y_i = 1 \mid \mathcal{Q}_{t-1}) \propto \gamma + \sum_{j \in \mathcal{Q}_{t-1},\, y_j = 1} w_{ij}, \tag{11}$$

where $w_{ij} = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2 / \sigma^2)$ represents spatial similarity weight, and $\gamma > 0$ is a smoothing prior that prevents degenerate estimates in sparsely observed regions.

Non-Myopic Policy. Unlike greedy or active learning approaches that prioritize immediate reward or local information gain, the Active Search formulation explicitly considers the full remaining inspection horizon. We adopt the Efficient Non-myopic Search (ENS) algorithm (Jiang et al., 2017), which approximates the optimal Bayesian policy for sequential discovery.

Formally, the optimal value function with budget b remaining satisfies the Bellman recursion:

$$V_b(\mathcal{D}_t) = \max_{i \in \mathcal{U}_t} \mathbb{E}\left[y_i + V_{b-1}(\mathcal{D}_t \cup \{(\mathbf{x}_i, y_i)\}) \mid \mathcal{D}_t \right], \tag{12}$$

where the expectation integrates over uncertain outcomes. Note that exact computation is intractable due to exponential growth of the search space with horizon length.

ENS addresses this challenge by approximating the future value using a branch-and-bound pruning strategy: for each candidate location, ENS computes an upper bound on achievable future utility assuming optimistic outcomes. Candidates whose upper bounds fall below the current best expected utility are pruned without expensive lookahead evaluation. This enables tractable non-myopic planning that sacrifices short-term gains when beneficial for long-term cumulative discovery.

5. Data

We use the dataset from the *RELand* study (Dulce Rubio et al., 2024), which focuses on humanitarian demining operations in Colombia—the country with the second-largest number of landmine victims globally (Prem et al., 2025). Since 1990, official Colombian statistics have recorded more than 40,000 landmine events resulting in over 12,000 casualties, 40% of them civilians (DAICMA, 2022), averaging more than three events and one victim daily for the past 35 years. This large-scale landmine contamination is a result

of its decades-long internal armed conflict, during which armed groups deployed landmines strategically to protect areas of interest including coca crops, drug trafficking routes, and illegal mining sites (BBC News, 2010). Notably, landmine contamination is highly concentrated geographically, with 50% of historical victims found in just 2% of the country's municipalities (Centro Nacional de Memoria Histórica & Fundación Prolongar, 2017).

The data covers 15 municipalities in the Antioquia region that have completed mine action operations and been declared mine-free. Each municipality is divided into 500 by 500 meter grid cells, resulting in 26,761 cells. Grid cell labels reflect the operational workflow of humanitarian demining organizations: a cell receives a positive label if it intersects a confirmed hazard area or cleared zone where mines were discovered during humanitarian demining operations. Otherwise, cells are labeled negative, indicating they were inspected and declared safe. This labeling scheme ensures that historical landmine events, which demining organizations use for planning surveys and interventions, can reliably predict contaminated areas. Moreover, the dataset includes 63 predictive features including geographic characteristics (e.g., altitude, soil texture, land use), socio-demographic variables (e.g., distance to settlements, schools, and roads), and remnants of war indicators (distance to coca crops and previous landmine reports) (Dulce Rubio et al., 2024).

The data exhibits significant class imbalance with only 1.6% positive cases (430 mines across 26,761 grid cells). Table 2 in the Appendix summarizes the dataset statistics for each municipality, showing substantial variation in both contamination density (0.03% to 10.22%) and operational scale (519 to 5,416 grid cells per municipality).

6. Results

We apply the six continuous learning strategies to each of the 15 municipalities with available ground truth labels. For each municipality, we allocate a query budget of 50% of grid cells and initialize Random Forest models using data from the other 14 municipalities as training data, emulating realistic geographic distribution shifts. Table 1 presents the final recall achieved by each method, showing clear performance differences between static, greedy, and adaptive approaches.

Across all municipalities, Greedy Retraining achieves the highest average recall (98.7%) at the 50% budget level, substantially outperforming static Greedy selection (84.9%), Explore-Exploit (88.4%), and Active Learning (74.5%). This performance gap demonstrates the central role of continuous model adaptation and fine-tuning for mine action operations, particularly when training and deployment regions differ substantially. San Francisco and San Roque municipalities illustrate this point: from the same random forest base model, static methods achieve low recall (22%), yet Greedy Retraining recovers to achieve perfect recall (100%) by rapidly correcting miscalibrated risk estimates based on field observations.

Notably, methods that explicitly balance exploration and exploitation underperform simpler Greedy Retraining despite their theoretical properties. This counterintuitive result reflects a fundamental objective mismatch: exploration-based strategies prioritize reducing model uncertainty to optimize final predictive performance when the labeling budget is exhausted, whereas mine action operations require maximizing immediate recall (mines found) during data acquisition. In this setting, exploring high-uncertainty regions often means surveying ambiguous negative instances rather than discovering missed positives, wasting limited query budget on locations that do not advance the operational goal.

Figure 1 demonstrates that this behavior persists consistently across budget levels from 0% to 50% of grid cells. Greedy Retraining dominates throughout the entire search process with the highest detection rate, approaching perfect recall as the 50% budget is exhausted. Explore-Exploit tracks closely behind, outperforming static Greedy Risk across most of the budget range, while Active Learning performs substantially worse than all greedy variants. This demonstrates that uncertainty-driven exploration offers limited value for recall-focused objectives in severely imbalanced settings where discovering rare positives, not reducing global model uncertainty, determines operational success.

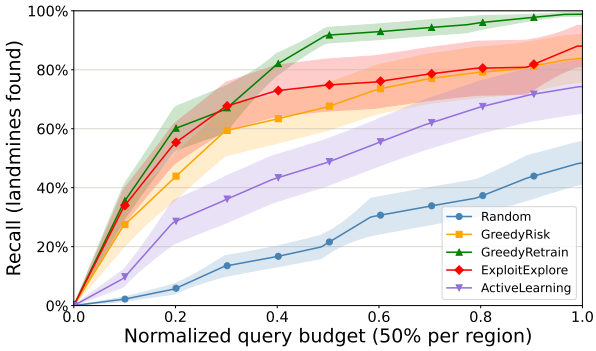

Figure 1. Average recall vs. search budget across municipalities ($\pm SD/\sqrt{n}$). Greedy Retraining consistently outperforms other methods for all budget levels. Experiments are run for T=10 rounds with batch size adjusted accordingly.

Active Search represents a fundamentally different scenario than the other methods, as it assumes no historical training data is available and must learn entirely from field observations. Despite this initial disadvantage, Active Search achieves 78.7% average recall, outperforming Active Learning (74.5%) that leverage training data from 14 other municipalities, demonstrating its potential for deployment in

Table 1. Final recall (%) at 50% budget across municipalities. Best method per municipality in bold. Active Search not completed for Sonsón due to computational constraints. Hyperparameters: $\alpha = 0.75$, $\beta = 0.5$, $w_t \in [0.6, 0.9]$, $\gamma = 0.1$.

Region	Mines	Budget	Random	Greedy	Explore-Exploit	Active Learning	Greedy Retraining	Active Search
Abejorral	11	1,025	45	**100**	91	91	91	90
Alejandría	5	259	0	80	60	**100**	**100**	60
Chigorodó	1	1,468	0	**100**	0	0	**100**	0
Cocorná	26	490	50	96	96	**100**	**100**	89
El Carmen	34	858	44	**100**	**100**	79	**100**	88
Granada	79	386	48	**99**	93	82	97	90
La Unión	1	339	**100**	**100**	**100**	**100**	**100**	**100**
Nariño	49	643	53	88	96	90	**100**	71
Sabanalarga	5	537	40	**100**	**100**	**100**	**100**	80
San Carlos	50	1,455	56	**98**	**98**	70	**98**	82
San Francisco	27	725	37	22	**100**	93	**100**	89
San Luis	16	864	50	**100**	**100**	38	**100**	87
San Rafael	38	739	50	**95**	**95**	87	**95**	76
San Roque	2	859	**100**	0	**100**	0	**100**	**100**
Sonsón	86	2,730	51	95	97	93	**100**	–
Average	28.7	891.8	48.2	84.9	88.4	74.5	**98.7**	78.7

settings where historical data is unavailable. However, computational constraints limit scalability as Active Search failed to complete on Sonsón, the largest municipality in the dataset, highlighting the tradeoff between non-myopic planning and operational feasibility.

Figure 2 illustrates Active Search behavior for San Francisco municipality at finer temporal resolution. Unlike methods initialized with pretrained models, Active Search exhibits an initially flat recall trajectory during an exploration phase of approximately 50 queries, where the algorithm builds its spatial contamination model from scratch. Once early targets are identified, recall increases sharply as the method exploits spatial clustering patterns and feedback structure, ultimately achieving 88.8% recall despite high sparsity (1.86% landmine density).

These results demonstrate that continuous model adaptation is crucial for effective landmine discovery under budget constraints and geographic distribution shift. When a reliable pretrained model is available, Greedy Retraining can deliver near-optimal performance with minimal algorithmic complexity, making it a practical choice for operational deployment. In data-scarce regimes, Active Search offers a compelling alternative achieving competitive performance despite learning from scratch, though computational constraints limit its applicability to smaller operational scales. The next section translates these empirical findings into concrete deployment recommendations in demining operations.

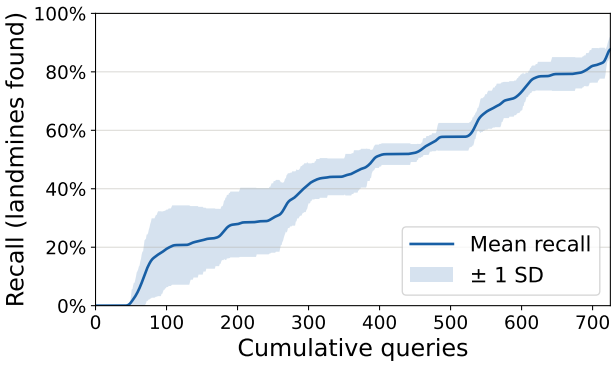

Figure 2. Recall vs. budget for San Francisco using Active Search ($\pm SD$; 10 runs). Initial flat phase represents exploration before rapid discovery, achieving 88.8% final recall.

7. Implications for Deployment

Our evaluation provides evidence of a missed opportunity in current AI-informed mine action systems: existing proposals operate statically when model adaptation could substantially improve performance. Continuous retraining approaches demonstrate considerable performance gains over static methods by incorporating field observations to refine risk estimates. This model fine-tuning proves particularly valuable under geographic distribution shift, where training data comes from different regions than deployment contexts. Importantly, these benefits can be achieved through straightforward retraining workflows and not necessarily involve sophisticated algorithmic frameworks.

Our results show that exploration-based methods, which balance uncertainty reduction with exploitation, are less effective than simpler greedy strategies in this domain. This finding is driven by a fundamental objective mismatch: exploration-focused approaches typically optimize for final model quality after budget exhaustion, whereas mine action operations prioritize maximizing discoveries during sequential querying. In our class imbalance setting, this exploration tends to select highly uncertain instances that may not optimize recall objectives. These patterns appear

consistent across budget levels, suggesting the findings may generalize beyond our specific evaluation settings.

7.1. Considerations for Method Selection

These empirical results suggest key considerations when using continuous learning methods in operational contexts:

When historical training data is available simple greedy retraining appears to offer substantial benefits. This approach shows strong performance when initial models could be trained on data from similar regions, even under moderate distribution shift, making it potentially feasible for field deployment with typical computational resources.

When training data is limited or distribution shift is severe active search methods offer value by learning directly from field observations without relying on pretrained models. Our results suggest this approach can achieve competitive performance when starting from minimal prior knowledge, though computational requirements may limit applicability to smaller operational scales. This strategy might be particularly relevant for initial deployments in new conflict contexts or when early surveys reveal severe model miscalibration. The strong performance of active search is driven by its explicit modeling of spatial clustering, where graph-based formulations that exploit these local dependencies may be better aligned with the underlying contamination structure than methods treating locations independently.

Exploration-based strategies are typically not suited for recall-focused settings. These methods allocate query budget to uncertainty reduction rather than positive discovery, potentially limiting effectiveness for mine action objectives.

7.2. Ongoing Deployment Efforts

These findings directly inform our ongoing collaboration with demining organizations, where we have deployed a pretrained risk assessment model in Colombia that operates under a static greedy approach. Field teams received an initial contamination risk map and prioritized survey locations based on these fixed risk scores. This system has already provided operational value in reducing both false alarms and clearance times, but feedback from field teams suggests significant distribution shift between the available training data (from other Colombian regions) and our deployment contexts.

Informed by our empirical findings, we are developing a user interface that enables retraining on demand, allowing field operators to trigger model updates as survey data becomes available. This implementation follows the greedy retraining paradigm validated in our empirical evaluation, where after each batch selection, the system generates updated risk predictions to guide subsequent survey batches. This ongoing deployment provides a critical test case for continuous learning in operational mine action under realistic conditions, where operational feedback will assess the practical feasibility of retraining workflows in mine action operations and measure performance improvements relative to the current static approach.

8. Conclusion

Our formalization and systematic evaluation of continuous learning strategies using operational data from Colombia demonstrates that AI-informed systems can substantially improve landmine discovery when models adapt to field observations. Simple greedy retraining achieves near-optimal performance with minimal computational requirements, while active search methods show promise when no training data is available, establishing continuous learning as a viable path forward for AI-guided humanitarian mine action operations. Integrating the non-myopic planning benefits of active search with pretrained models in computationally feasible ways remains an important direction for future work. These adaptive AI systems have the potential to accelerate land clearance and save lives in conflict-affected communities, with deployment efforts already underway to translate these findings into operational practice.

Acknowledgments

This research was supported by the Minesight Seedling Program. The opinions expressed are those of the authors and do not reflect the views of Minesight. The authors declare no conflict of interest. All errors remain our own.

References

BBC News. How landmines complicate Colombia's drugs fight. https://www.bbc.com/news/world-latin-america-11980034, 2010. Accessed: 2023-05-20.

Centro Nacional de Memoria Histórica and Fundación Prolongar. *La guerra escondida: Minas Antipersonal y Remanentes Explosivos en Colombia*. Centro Nacional de Memoria Histórica, Bogotá, 2017.

Cirillo, F., Solmaz, G., Peng, Y.-H., Bizer, C., and Jebens, M. Desk-aid: Humanitarian aid desk assessment with geospatial ai for predicting landmine areas. *arXiv preprint arXiv:2405.09444*, 2024.

Collins, R., Fragniere, L., and Dulce Rubio, M. Advancements in mine action: Enhancing remote reporting and analysis through innovative technologies. *The Journal of Conventional Weapons Destruction*, 28(3):7, 2024.

DAICMA. Situación víctimas minas antipersonal en colombia: Datos abiertos colom-

bia. https://www.datos.gov.co/Inclusi-n-Social-y-Reconciliaci-n/Situaci-n-V-ctimas-Minas-Antipersonal-en-Colombia/yhxn-eqqw/about_data, March 2022. Accessed: 2024-07-13.

Dulce Rubio, M., Zeng, S., Wang, Q., Alvarado, D., Moreno Rivera, F., Heidari, H., and Fang, F. RELand: Risk estimation of landmines via interpretable invariant risk minimization. *ACM Journal on Computing and Sustainable Societies*, 2(2):1–29, 2024.

Garnett, R., Krishnamurthy, Y., Xiong, X., Schneider, J., and Mann, R. Bayesian optimal active search and surveying. In *Proceedings of the 29th International Conference on International Conference on Machine Learning*, pp. 843–850, 2012.

GICHD. A guide to international mine action standards 2010. https://www.gichd.org/publications-resources/publications/a-guide-to-the-international-mine-action-standards-2010/, October 2009. Accessed: 2023-02-28.

Jiang, S., Malkomes, G., Converse, G., Shofner, A., Moseley, B., and Garnett, R. Efficient nonmyopic active search. In *International Conference on Machine Learning*, pp. 1714–1723. PMLR, 2017.

Jiang, S., Malkomes, G., Abbott, M., Moseley, B., and Garnett, R. Efficient nonmyopic batch active search. *Advances in Neural Information Processing Systems*, 31, 2018.

Kischelewski, B., Cathcart, G., Wahl, D., and Guedj, B. Ai for erw detection in clearance operations–the state of research. *arXiv preprint arXiv:2411.05813*, 2024.

Prem, M., Purroy, M. E., and Vargas, J. F. Landmines: The local effects of demining. *Journal of Public Economics*, 247:105399, 2025.

Sarkar, A., DiChristofano, A., Das, S., Fowler, P. J., Jacobs, N., and Vorobeychik, Y. Active geospatial search for efficient tenant eviction outreach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 28340–28348, 2025.

Settles, B. Active learning literature survey. Technical Report 1648, University of Wisconsin–Madison, January 2009.

UNMAS. The deadly legacy of landmines. https://news.un.org/en/story/2023/04/1135252, 2023. Accessed: 2024-07-13.

A. Appendix

A.1. Dataset Description

We use the RELand dataset from Dulce Rubio et al. (2024), covering 15 municipalities in the Antioquia region of Colombia, comprising 26,761 grid cells of 500 by 500 meters each.[1] The dataset includes 63 predictive features across three categories: (i) Geographic characteristics: altitude, animal density, soil texture, temperature, rainfall, distance to rivers, land use, weather, and topographic relief; (ii) Socio-demographic variables: distance to airports, seaports, settlements, financial institutions, schools, buildings, and roads, relative wealth index , population, and night lights index; and (iii) Remnants of war indicators: distance to coca crops, deforestation, mining, and power, telecommunication and oil lines. Overall contamination density, measured as the percentage of grid cells containing mines, is 1.6%, reflecting significant class imbalance. Table 2 breaks down contamination rates and operational scale by municipality, showing substantial variation across the region.

Table 2. Dataset statistics for finalized operations in 15 municipalities in Antioquia, Colombia.

Municipality	**Grid cells**	**Mines**	**Density (%)**
Abejorral	2,051	11	0.54
Alejandría	519	5	0.96
Chigorodó	2,937	1	0.03
Cocorná	980	26	2.65
El Carmen	1,716	34	1.98
Granada	773	79	10.22
La Unión	678	1	0.15
Nariño	1,287	49	3.81
Sabanalarga	1,074	5	0.47
San Carlos	2,911	50	1.72
San Francisco	1,450	27	1.86
San Luis	1,728	16	0.93
San Rafael	1,478	38	2.57
San Roque	1,718	2	0.12
Sonsón	5,461	86	1.57

A.2. Implementation Details

All experiments are run for $T = 10$ rounds, with batch size b_t adjusted according to the maximum budget in each municipality, set to 50% of the total grid cells. Table 3 summarizes the key hyperparameters used.

Table 3. Hyperparameter settings used in our experiments.

Parameter	Definition	Value
α	Exploitation fraction in explore-exploit strategy	0.75
β	Active learning uncertainty-density tradeoff parameter	0.5
w_t	Active learning exploitation weight (increases linearly with remaining budget)	[0.6, 0.9]
γ	Active search smoothing prior	0.1

[1]The data is publicly available at: `https://huggingface.co/datasets/mdulcer10/landmine-detection`.