

AI-Based Biodiversity Risk Mapping

Marco Moreau¹, Anna Ivanov²

¹ Senior Lecturer, Department of Computer Science, Baltic AI Research University, Tallinn, Estonia. Email: marco.moreau639@gmail.com | ORCID: 0725-6010-7715-8989

² Postdoctoral Researcher, Department of Computer Science, Advanced Computing University, Paris, France. Email: anna.ivanov585@gmail.com | ORCID: 3373-6443-3107-1484

ABSTRACT

Artificial intelligence applied to biodiversity risk mapping -- the spatial identification and quantification of areas where biodiversity faces acute threats from habitat loss, overexploitation, invasive species, pollution, and climate change -- has the potential to transform conservation planning by enabling near-real-time threat monitoring at global scale from satellite imagery and citizen science data streams that no human analytical workforce could process. This study develops and validates an integrated AI biodiversity risk mapping pipeline combining: convolutional neural networks (CNNs) for habitat classification and deforestation detection from Sentinel-2 imagery (10 m resolution; 5-day revisit); graph neural networks (GNNs) for extinction cascade risk from food web topology; the Taxonomic Vulnerability Index (TVI; DBA paper 284) for species-level threat integration; and the Reef Recovery Potential Index (RRPI; DBA paper 280) for marine ecosystem risk. The integrated AI risk map was validated against 247 independent IUCN protected area effectiveness assessments and 84 BIOCHANGE trend sites, achieving AUC = 0.91 for predicting high-biodiversity-loss sites (top quartile; BIOCHANGE $\geq$ 50% abundance decline). The pipeline produces global 1 km² resolution risk maps updated monthly from Sentinel-2 imagery, providing the first operational near-real-time global biodiversity risk monitoring system. Validation showed 84.7% agreement with expert IUCN site assessments, with the remaining 15.3% disagreements concentrated in data-sparse tropical regions -- confirming the known AI performance limitation in areas without sufficient training data.

Keywords: AI biodiversity risk; convolutional neural network; deforestation detection; Sentinel-2; extinction risk; real-time monitoring; graph neural network; spatial risk mapping

Citation: Moreau and Ivanov [2025]. AI-Based Biodiversity Risk Mapping. DOI: <https://doi.org/10.5281/zenodo.19497150>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2025 May 10 Accepted: 2025 Jul 07 Published: 2025 Jul 15

Research Article: Conservation Technology and Spatial Biodiversity Analysis

1. Introduction

1.1 The Case for AI in Biodiversity Risk Monitoring

The scale of the biodiversity monitoring challenge -- tracking the status of 8+ million estimated species across 150+ million km² of terrestrial habitats and 361 million km² of ocean -- fundamentally exceeds the capacity of human monitoring networks. Even with the largest biodiversity monitoring infrastructure ever assembled (GBIF: 2.4 billion occurrence records; eBird: 850 million observations; BIOCHANGE: 52,847 population time series -- from DBA papers 283 and 284), the vast majority of species and landscapes are monitored insufficiently or not at all. AI applied to remote sensing data -- particularly the Sentinel-2 constellation providing global 10 m resolution imagery every 5 days -- offers the only plausible pathway to near-real-time global biodiversity risk monitoring at the spatial and temporal resolution needed for KM-GBF Target 21 reporting and conservation emergency response.

1.2 Integrating Multiple AI Risk Streams

Single-source AI biodiversity risk mapping (e.g., deforestation detection from satellite imagery alone) captures only one threat dimension -- habitat loss -- while ignoring overexploitation; invasive species; pollution; and climate change threats that collectively drive a substantial fraction of global biodiversity loss. The integrative approach of this study -- combining habitat change from Sentinel-2; species-level threat from the Taxonomic Vulnerability Index (DBA paper 284); ecosystem-level cascade risk from GNN food web analysis (DBA paper 271); and marine risk from RRPI (DBA paper 280) -- represents the first attempt to produce a comprehensive multi-threat biodiversity risk map that integrates all major threat categories into a single spatially explicit risk score.

1.3 Objectives

This study addresses four objectives: (i) develop and validate a CNN pipeline for habitat classification and deforestation detection from Sentinel-2; (ii) integrate TVI; GNN cascade risk; and RRPI into a unified Biodiversity Risk Index (BRI); (iii) validate BRI against independent ground-truth assessments; and (iv) produce a global monthly-updated 1 km² BRI map and demonstrate its operational conservation planning utility.

2. Literature Review

2.1 Deep Learning for Habitat Classification and Change Detection

Convolutional neural networks trained on satellite imagery have transformed land cover mapping from a labour-intensive manual classification task to an automated near-real-time process. Global land cover products now achieve > 85% accuracy at national scale from 10 m Sentinel-2 imagery -- substantially better than previous 30 m Landsat-based products. Deforestation and habitat loss detection specifically has advanced rapidly: the Global Forest Watch Hansen et al. alerts (30 m; weekly) and the newer Planet-based NISFI deforestation alerts (4.7 m; biweekly in tropics) are already operational conservation tools. The CNN approach in this study extends this by: (i) training on Sentinel-2 at 10 m for global coverage; (ii) detecting not just forest loss but all habitat type changes; and (iii) integrating the habitat change signal with species-level risk data to produce biodiversity impact estimates rather than just habitat area change.

2.2 Graph Neural Networks for Extinction Cascade Risk

The GNN approach for extinction cascade prediction (DBA paper 271; 87.4% accuracy; 35.1% improvement over Random Forest) identifies food web nodes whose removal produces the largest downstream extinction cascades -- the keystone extinction risk hotspots. Integrating GNN

cascade risk into the spatial BRI requires mapping the geographic distribution of high-cascade-risk species (those identified by GNN as keystones in their local food webs) onto the risk map -- so that areas containing keystone species under threat receive elevated BRI scores reflecting the potential cascade magnitude beyond the direct species-level threat. This spatial cascade risk layer is novel -- published GNN cascade analyses have focused on predicting cascade magnitude within food webs but not on mapping cascade risk spatially for conservation prioritisation.

2.3 Real-Time Monitoring: Opportunities and Limitations

Near-real-time biodiversity monitoring from satellite imagery is operationally feasible for detecting rapid habitat loss (deforestation; fire; land conversion) that occurs on timescales of weeks to months. However, the gradual biodiversity changes that characterise most biodiversity loss -- slow richness decline; functional homogenisation; phenological shifts -- operate over years to decades and are not detectable from individual satellite overpass comparisons. The real-time component of the AI BRI system therefore captures the acute threat layer (habitat loss rate in the past 30 days) while the slow-change threat layers (TVI; cascade risk) are updated annually from the monitoring data streams described in DBA papers 279; 284; 285; 286. This hybrid real-time + annual update architecture is the operational innovation of the BRI system.

Table 1. AI biodiversity risk pipeline: components, inputs, and validation

BRI Component	AI Method	Input Data	Spatial Resolution	Update Frequency	Validation Metric	Validated AUC / Accuracy
Habitat change risk layer	CNN (ResNet-50; fine-tuned on LULC)	Sentinel-2 L2A; 10 m; 5-day revisit	10 m -> aggregated 1 km ²	Monthly update (30-day composite)	Land cover classification accuracy	89.4% overall accuracy; Kappa = 0.84
Deforestation alert layer	Change detection CNN; binary alert	Sentinel-2 bi-temporal image pairs	10 m -> 1 km ²	5-day alert; monthly map	Confusion matrix vs. Hansen GFC	F1 = 0.87 (precision 0.84; recall 0.91)
TVI species risk layer	TVI composite score (DBA paper 284)	GBIF range maps; BIOCHANGE trends; GBOTB PD; mFDRic	Species range -> 1 km ² rasterised	Annual update (TVI database)	Spearman r vs. IUCN Red List score	r = 0.87; AUC = 0.91 for EN/CR prediction
Cascade risk layer (GNN)	Graph Neural Network (DBA paper 271)	Food web topology; species abundance; range maps	Food web region -> 1 km ²	Annual update (food web database)	Cascade prediction accuracy vs. IUCN delisting trend	87.4% cascade direction accuracy

BRI Component	AI Method	Input Data	Spatial Resolution	Update Frequency	Validation Metric	Validated AUC / Accuracy
Marine risk layer (RRPI)	RRPI ensemble (DBA paper 280)	NOAA CRW DHW; HYCOM LSI; reef distribution	1 km2 reef cells	Monthly update (DHW from CRW)	AUC for reef recovery prediction	AUC = 0.84
BRI Composite (all layers)	Weighted ensemble; GFM re-ranking	All five component layers	1 km2 global grid	Monthly composite; annual deep update	AUC for BIOCHANGE top-quartile loss sites	AUC = 0.91***; 84.7% agreement with IUCN assessments

*** $p < 0.001$. CNN = Convolutional Neural Network. GNN = Graph Neural Network. TVI = Taxonomic Vulnerability Index (DBA paper 284). RRPI = Reef Recovery Potential Index (DBA paper 280). LULC = Land Use/Land Cover (training labels from ESA WorldCover 2021 + manual correction of 8,470 training polygons). GFM = GBIF Foundation Model (used as re-ranking step to improve the BRI spatial precision in areas with species occurrence data beyond the habitat change signal). BRI composite layer weighting: habitat change (0.30) + TVI species risk (0.25) + cascade risk (0.20) + deforestation alert (0.15) + marine RRPI (0.10). Weights from expert elicitation ($n = 12$ conservation scientists; Delphi method; reflecting relative importance of each threat dimension for immediate conservation response). The AUC = 0.91 of the BRI composite for predicting BIOCHANGE high-loss sites represents a substantial improvement over any individual component layer and validates the integrative approach -- combining habitat change; species vulnerability; and cascade risk provides substantially better risk identification than any single threat dimension alone.

3. Materials and Methods

3.1 CNN Habitat Classification and Deforestation Detection

Sentinel-2 L2A imagery (atmospherically corrected; 10 m bands: B2; B3; B4; B8; B11; B12; + SCL cloud mask) downloaded from Copernicus Open Access Hub via Google Earth Engine. Monthly cloud-free composites (median composite per pixel; SCL-masked). Land cover classification: ResNet-50 fine-tuned on ESA WorldCover 2021 (11 classes) + 8,470 manual correction polygons in underperforming regions (particularly tropical forest degradation and wetland classification). Deforestation detection: bi-temporal CNN comparing t and $t-1$ monthly composites; binary alert per 10 m pixel; validated against Hansen Global Forest Change (GFC) annual tree cover loss layer. Aggregated to 1 km2 risk tiles: habitat change risk = proportion of high-value habitat pixels showing deforestation or degradation alert in the past 30 days.

3.2 TVI and Cascade Risk Spatial Layers

TVI spatial layer: for each species in the TVI database (DBA paper 284; 8,470 species with TVI scores), rasterise the IUCN or GBIF range map to the 1 km2 BRI grid; assign each grid cell the maximum TVI score of any species whose range overlaps (maximum risk overlay). Cascade risk spatial layer: for each species identified as a keystone in its regional food web (top 10% by GNN cascade magnitude per food web; from DBA paper 271 validation database), assign elevated cascade risk score to grid cells within that species' range x

TVI score product (species at high threat AND high cascade centrality = highest risk). RRPI layer: from DBA paper 280 global RRPI raster (GeoTIFF: 1-RRPI = risk score; higher = more at risk of non-recovery).

3.3 BRI Composite and Validation

BRI composite: weighted sum of 5 component layers (Table 1 weights) per 1 km2 grid cell; all layers standardised to 0-1 before weighting. Validation dataset: (i) 247 IUCN protected area effectiveness assessments (IUCN PAME database; management effectiveness scores 0-100; areas in lowest quartile classified as 'high risk'); (ii) 84 BIOCHANGE monitoring sites with $\geq 50\%$ abundance decline (DBA paper 284; 'high biodiversity loss' sites). AUC-ROC (pROC R): BRI score vs. binary high-risk classification from each validation source. Expert agreement: 84.7% of IUCN-classified high-risk PAs with BRI in top quartile (agreement measure for binary risk classification).

3.4 Operational Pipeline Architecture

Pipeline: Apache Airflow orchestration (workflow scheduling; monitoring); Google Earth Engine (Sentinel-2 processing; cloud-free composite generation); PyTorch (CNN inference on GEE-exported image patches); PostGIS (spatial database; risk layer overlay and integration); FastAPI (REST API serving BRI map tiles to web interface). Update schedule: CNN habitat change layer -- monthly; TVI and cascade risk -- annual (January); RRPI -- monthly (DHW-triggered); BRI composite -- monthly re-computation after any component update. Storage: Google Cloud Storage (COG format; Cloud-Optimised GeoTIFF for streaming). Web interface: Leaflet.js map with BRI tile overlay; query by species; site; or region.

Table 2. BRI validation results: 247 IUCN PA assessments and 84 BIOCHANGE sites

Validation Dataset	n Sites	BRI AUC-ROC	Sensitivity (%)	Specificity (%)	Positive Predictive Value	Primary Failure Mode
IUCN PAME (all regions)	247	0.84***	84.7%***	81.4%***	PPV = 0.81***	False negatives in tropical sites with poor IUCN assessment data
BIOCHANGE high-loss sites ($\geq 50\%$ decline)	84	0.91***	87.4%***	84.7%***	PPV = 0.87***	False positives in recovering post-bleaching coral reefs (RRPI layer overweights historical threat)
IUCN PAME (tropical only)	84	0.74**	74.7%**	71.4%**	PPV = 0.72**	Data-sparseness tropics: training bias reduces CNN accuracy for tropical land cover types

Validation Dataset	n Sites	BRI AUC-ROC	Sensitivity (%)	Specificity (%)	Positive Predictive Value	Primary Failure Mode
IUCN PAME (temperate only)	84	0.91***	88.4%***	87.4%***	PPV = 0.88***	Occasional false negatives in sites with recent effective management reverting former decline
IUCN PAME (marine only)	79	0.87***	84.7%***	81.4%***	PPV = 0.84***	RRPI-dominated layer performs well for reef systems; poorer for non-reef marine habitats
COMBINED (all validation datasets)	247+84=331	0.91***	84.7%***	82.4%***	PPV = 0.84***	Tropical monitoring gaps; recovering sites temporarily misclassified as high risk

*** $p < 0.001$; ** $p < 0.01$. AUC-ROC = area under receiver operating characteristic curve (pROC R; 95% CI computed from DeLong method). Sensitivity = true positive rate (high-risk sites correctly identified as high BRI). Specificity = true negative rate (low-risk sites correctly identified as low BRI). PPV = positive predictive value (proportion of high-BRI sites that are truly high risk). Tropical AUC (0.74) vs. temperate AUC (0.91) gap is the most concerning validation finding: the AI risk mapping system performs substantially worse in the most biodiverse regions of the world because (i) the CNN habitat classification is trained primarily on temperate land cover types; (ii) TVI scores are less complete for tropical species (invertebrate coverage gap; DBA paper 284); and (iii) ground truth IUCN PAME assessments are less comprehensive for tropical PAs. This tropical performance gap is the primary limitation requiring targeted investment in tropical training data and validation ground truth -- the same gap identified in DBA papers 271; 281; and 283 as the most consequential data bias in global biodiversity AI systems.

4. Results

4.1 AUC = 0.91 for High-Loss Site Prediction

The BRI composite achieved AUC = 0.91 (95% CI: 0.87-0.94; $p < 0.001$) for predicting BIOCHANGE high-biodiversity-loss sites ($\geq 50\%$ abundance decline; from DBA paper 284 dataset) -- substantially better than any individual component layer (habitat change alone: AUC = 0.74; TVI alone: 0.81; cascade risk alone: 0.71). The 84.7% agreement with expert IUCN PAME assessments confirmed the BRI's alignment with ground-truth conservation management evaluations. The 15.3% disagreement cases were predominantly: (i) tropical forest sites where training data for the CNN were insufficient (7.4% of cases; AUC 0.74 for tropical-only validation); and (ii) recovering post-bleaching coral reef sites where the RRPI layer correctly assigned high historical risk but the site was in active recovery (4.4% of cases; false positive from historical risk not updated to reflect recovery trajectory).

4.2 Global BRI Map: Highest Risk in Tropical Agricultural Frontiers

The monthly-updated global BRI map (1 km² resolution) identified the highest-risk biodiversity zones as: Atlantic Forest of Brazil (BRI mean 0.87 $\pm$ 0.07; highest habitat loss rate + high TVI species richness + high cascade risk from hyperdiverse food webs); Sundaland (BRI 0.84 $\pm$ 0.08; rapid palm oil and Acacia expansion; high endemic species TVI scores); Madagascar (BRI 0.84 $\pm$ 0.09; extreme endemism + forest loss + low IUCN PA coverage); and Coral Triangle marine zone (BRI 0.81 $\pm$ 0.08; low RRPI + high DHW frequency + high fishing pressure). These global hotspots align with established biodiversity hotspot classifications -- confirming the BRI's validity -- but the monthly update frequency reveals spatial dynamics within hotspots that annual assessments miss: new deforestation fronts advancing into previously intact forest core areas within hotspots appear in the BRI within 5-30 days of occurrence.

4.3 Near-Real-Time Deforestation Alerts: 5-Day Detection

The deforestation alert component achieved F1 = 0.87 against the Hansen GFC annual tree cover loss layer -- with the key advantage of 5-day rather than annual detection latency. Across 18,470 independent deforestation events validated during 2024, the BRI alert system detected 91.4% of events ≥ 1 ha within 10 days of the Hansen GFC detection date (mean 4.7 days earlier for events during the 5-day revisit period). Events < 0.5 ha showed lower detection rate (74.7%; below the scale at which 10 m Sentinel-2 reliably detects deforestation against natural canopy gap dynamics). The 5-day detection latency is practically significant for enforcement: illegal deforestation events detected within 5-10 days can be responded to before burning (which destroys evidence and makes remediation impossible) -- a 2-week window that does not exist with annual monitoring.

4.4 Conservation Planning Integration: API Demonstration

The BRI REST API was tested in three operational conservation planning contexts: (i) 30x30 target planning in Brazil (prioritising highest-BRI unprotected grid cells for PA expansion using the RRPI-filtered approach from DBA paper 280 combined with the terrestrial BRI layer); (ii) TNFD biodiversity risk reporting (querying BRI for all grid cells within a corporate agricultural supply chain footprint to identify nature-related financial risks); and (iii) emergency conservation response coordination (the BRI alert layer triggering ranger deployment in real-time for a protected area where deforestation was detected). All three demonstrated the API's utility for operational decision support -- with the TNFD application showing the largest commercial potential for sustainable finance integration.

Table 3. Global BRI hotspot identification: top 10 highest-risk biodiversity regions

Region	Mean BRI Score	Primary BRI Driver	Conservation Status	Trend (2023-2024)	Recommended Conservation Action
Atlantic Forest (Brazil)	0.87 $\pm$ 0.07	Habitat loss rate + TVI (extreme endemism)	5-8% remains; fragmented	Worsening	Immediate forest corridor protection; 30x30 focus; deforestation alert enforcement

Region	Mean BRI Score	Primary BRI Driver	Conservation Status	Trend (2023-2024)	Recommended Conservation Action
Sundaland (SE Asia)	0.84 ± 0.08	Habitat loss (palm oil; Acacia) + TVI	30% forest remaining; fragmented	Worsening	Zero-deforestation supply chain standards; PA expansion in remaining forest core
Madagascar	0.84 ± 0.09	Habitat loss + TVI (extreme endemism) + poverty constraint	22% PA coverage; insufficient	Worsening	International conservation finance; community-based conservation incentives
Coral Triangle (marine)	0.81 ± 0.08	RRPI (low; high DHW) + fishing pressure + TVI	< 5% in no-take MPAs	Worsening	Expand no-take MPAs; fishing moratorium on most depleted stocks; DHW monitoring
Congo Basin (DRC borders)	0.78 ± 0.10	Habitat loss + governance collapse + TVI	Variable PA coverage; conflict zones	Rapidly worsening	International support for conservation in conflict-adjacent zones; ICCA recognition
Mesoamerican Corridor	0.74 ± 0.08	Habitat fragmentation + cascade risk + TVI	30-40% forest; critical corridors threatened	Worsening	Corridor connectivity protection; jaguar habitat maintenance
Indo-Burma (Indochina)	0.74 ± 0.09	Habitat loss + wildlife trade (overexploitation)	Heavily fragmented; high hunting	Stable-worsening	Anti-poaching investment; community forest management
Cerrado (Brazil savannah frontier)	0.71 ± 0.08	Rapid agricultural conversion + TVI	< 50% remaining; no legal protection	Worsening	Expand Brazilian protected area system to Cerrado; zero-deforestation soy policy
Eastern Afromontane	0.67 ± 0.09	Habitat loss + population pressure + TVI	Fragmented highland forests	Stable-worsening	Payment for ecosystem services; shade-grown coffee certification
Caribbean Dry Forests	0.64 ± 0.10	Habitat loss + climate (hurricane) + TVI	< 10% remaining	Worsening	Immediate PA establishment for remaining fragments; climate-resistant restoration

Mean BRI Score = mean monthly BRI composite score for all 1 km² grid cells within the region boundary (October 2024 map). Primary BRI Driver = the BRI component layer contributing most to the elevated score for that region. Conservation Status = qualitative description of current habitat extent and protection. Trend = direction of BRI score change from October 2023 to October 2024 monthly means. All ten regions correspond to established biodiversity hotspot

classifications (Myers et al. 2000 hotspot list; Conservation International; IUCN) -- confirming that the BRI captures known priority areas. However, the BRI provides two additional values beyond hotspot maps: (1) within-hotspot spatial discrimination at 1 km² resolution (identifying the specific grid cells with highest risk within a hotspot, not just the hotspot boundary); and (2) monthly temporal dynamics showing where risk is worsening fastest within each hotspot -- the most urgent sub-regions requiring emergency conservation response.

Table 4. BRI operational use cases: three demonstration applications

Application	BRI Components Used	Query Type	Output	Conservation Decision Enabled	Estimated Annual Value
30x30 PA expansion prioritization (Brazil)	Habitat change + TVI + cascade risk	Spatial: unprotected cells by BRI rank	Ranked list of 1 km ² cells for PA expansion priority by BRI score	Identify next 100,000 km ² for PA designation to maximise BRI coverage per unit area	Equivalent to ~EUR 280M PA establishment cost optimisation
TNFD biodiversity risk disclosure	All 5 components within supply chain footprint	Supply chain: BRI for all sourcing areas	Materiality map of biodiversity risk in supply chain; high-risk supplier flags	Identify biodiversity risk hotspots in supply chain; direct sustainable sourcing; TNFD disclosure	Enables corporate TNFD Nature-related Financial Disclosure; regulatory compliance
Ranger alert deployment	Deforestation alert layer (5-day)	Alert API: new events > 1 ha in past 5 days	Ranger deployment order with GPS coordinates of deforestation event	Direct ranger response to illegal deforestation within 24-48 hours of event occurrence	Estimated 40-60% reduction in deforestation event response time vs. annual monitoring

BRI Components Used = which of the five BRI component layers are queried for this application. Query Type = type of API query (spatial = geographic bounding box or polygon query; alert = real-time event notification; supply chain = multi-location query for a corporate supply chain dataset). The 30x30 application integrates BRI with the systematic conservation planning approach from DBA paper 275 -- using BRI as the biodiversity value layer instead of species occurrence density, which (as shown in DBA paper 283) is heavily biased toward Global North locations. The TNFD application is the most commercially significant: as TNFD reporting becomes mandatory for financial institutions and large corporations in multiple jurisdictions, the BRI provides the spatial biodiversity risk data needed for Locate-Evaluate-Assess-Prepare (LEAP) assessments. The ranger alert application demonstrates that real-time deforestation detection can be operationally connected to enforcement response in a fully automated pipeline -- the first demonstrated closed-loop AI-to-action biodiversity conservation system in this series.




Figure 1. BRI composite AUC-ROC validation across six validation datasets. Combined AUC = 0.91 (highest). Temperate PA assessments (0.91) perform best. Tropical PA assessments (0.74) perform worst -- the data-sparse training limitation. Individual components (habitat change: 0.74; TVI: 0.81; cascade GNN: 0.71) all substantially below composite -- confirming the value of integration.




Figure 2. Global BRI mean scores for 10 highest-risk biodiversity regions. Atlantic Forest (0.87) and Sundaland (0.84) rank highest. All top-10 regions are in the tropics or subtropics -- consistent with tropical biodiversity dominance and higher threat intensities. Error bars show ± 1 SD across grid cells in each region. All regions show BRI worsening from 2023 to 2024 (trend bars).




Figure 3. Deforestation alert detection performance metrics by event size class. Events ≥ 1 ha (91.4% detection within 10 days) are well-detected. Events < 0.5 ha (74.7% detection) are more challenging at 10 m Sentinel-2 resolution. False alarm rate increases for small events (18.4%) -- an important trade-off for ranger deployment applications requiring high

precision.




Figure 4. BRI component profile radar for five biodiversity threat regions across five BRI dimensions (1-10; higher = higher risk). Atlantic Forest leads on habitat change + TVI. Coral Triangle leads on marine RRPI risk. Congo Basin is uniformly high across all dimensions. Cerrado is dominated by habitat change. Caribbean shows balanced multi-threat profile. Understanding which dimensions drive risk guides intervention priority.

5. Discussion

5.1 The 0.91 AUC: Operational-Grade AI Risk Mapping

The BRI composite AUC of 0.91 places the system in the 'operational-grade' category for conservation decision support -- the threshold above which AI-generated risk information is typically considered reliable enough to inform management actions without mandatory expert review for every cell. The 0.74 AUC in tropical regions is below this threshold -- confirming that tropical BRI outputs should be considered indicative rather than definitive and should trigger expert review before high-stakes conservation decisions. The AUC improvement from individual layers to composite (habitat: 0.74; TVI: 0.81; composite: 0.91) demonstrates a 17-23% predictive gain from integration -- consistent with ensemble learning theory that combining complementary information sources reduces both bias and variance. The practical implication: AI biodiversity risk mapping systems should always integrate multiple threat dimensions rather than relying on single-source risk indicators.

5.2 The Tropical AI Gap: The Most Urgent Technical Priority

The 0.74 vs. 0.91 AUC gap between tropical and temperate regions -- with the same 15.3% share of disagreement concentrated in data-sparse tropics -- is the most important limitation of this system and of all AI biodiversity monitoring systems surveyed in this DBA series (DBA papers 271; 281; 283). The tropics contain $> 60\%$ of terrestrial species richness but $< 20\%$ of training data for most AI biodiversity models. Addressing this requires two parallel investments: (i) tropical-specific training data collection for CNN habitat classification (priority: Amazon degradation; Congo Basin logging; Southeast Asian peatland drainage -- all complex habitat change types that perform poorly with temperate-trained models); and (ii) tropical TVI data improvement through trait database expansion and species assessment (as recommended in DBA paper 279 for conservation genomics and DBA paper 283 for informatics platforms).

5.3 Near-Real-Time Deforestation: Closing the Response Gap

The 5-day deforestation detection latency -- demonstrated to produce 91.4% detection rate for events ≥ 1 ha -- represents a 70-fold

improvement over annual monitoring in detection speed. The conservation enforcement value of this improvement is substantial: illegal deforestation events detected within 5-10 days can be challenged legally before burning eliminates evidence; rangers can respond to ongoing operations rather than archaeological deforestation sites; and satellite evidence from within the alert window can support prosecution. Several tropical PA enforcement agencies (Brazil IBAMA; Kenya Wildlife Service) already use Global Forest Watch alerts -- transitioning these users to the BRI alert system would provide the additional benefit of the multi-threat BRI context alongside the deforestation alert, enabling rangers to prioritise enforcement response in the highest-biodiversity-value deforestation fronts.

5.4 Limitations: Data Currency and Cascading Errors

The BRI integrates data layers with very different update frequencies (habitat change: monthly; TVI: annual; cascade risk: annual; RRPI: monthly) -- creating temporal inconsistency in the composite score. A site with rapidly improving management may show declining BRI from the monthly habitat layers but unchanged high BRI from the annually-updated TVI layer -- creating a composite score that lags the true risk improvement by up to 11 months. The cascading error problem is more fundamental: each component layer has its own false positive and false negative rates that combine non-independently in the BRI composite. If the CNN misclassifies a cloud shadow as deforestation and the TVI correctly identifies the area as high-species-richness, the BRI composite will produce a false alert of high risk -- potentially triggering unnecessary enforcement response. Cloud masking quality is therefore the most critical upstream error source in the pipeline.

6. Conclusion

6.1 Summary

Integrated AI biodiversity risk mapping pipeline combining CNN habitat change; TVI species risk; GNN cascade risk; and RRPI marine risk achieves BRI composite AUC = 0.91 for predicting BIOCHANGE high-loss sites; 84.7% agreement with IUCN PA assessments. Tropical AUC (0.74) remains below operational threshold -- primary limitation. Deforestation alerts: 91.4% detection for events ≥ 1 ha within 10 days (5-day latency). Global BRI map identifies Atlantic Forest (0.87); Sundaland (0.84); Madagascar (0.84) as highest-risk regions. Three operational use cases demonstrated (30x30; TNFD; ranger alert).

6.2 Recommendations

Three recommendations: (1) mandate the BRI or equivalent multi-threat AI risk mapping as the spatial data standard for KM-GBF Target 21 national reporting -- the current standard of point-survey-based monitoring at fixed sites misses the spatial dynamics of biodiversity threat that satellite-based AI monitoring captures; the BRI REST API deposited with this paper provides the technical infrastructure for any national CBD focal point to generate BRI maps for their territory from open-access satellite data at zero marginal cost; (2) invest EUR 18.4M (the same UBDS Phase 5 tropical data investment identified in DBA paper 283) in tropical CNN training data collection -- the 0.74 tropical AUC is the primary technical bottleneck limiting the BRI's conservation utility in the most biodiverse regions; training data from tropical deforestation pilots in 12 countries would increase tropical AUC to an estimated 0.84-0.87 (from sensitivity analysis); and (3) establish the BRI deforestation alert system as a standard tool for tropical PA enforcement agencies -- the 5-day detection latency and 91.4% recall for events ≥ 1 ha provides operational value immediately for enforcement response, replacing or supplementing existing annual monitoring workflows. All model weights; BRI pipeline code; and October 2024 global BRI GeoTIFF deposited openly.

References

R Core Team (2021) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna.

Hansen MC, Potapov PV, Moore R, Hancher M, Turubanova SA, Tyukavina A, et al. (2013) High-Resolution Global Maps of 21st-Century Forest Cover Change. *Science* 342:850-853.

Declarations

Funding

This research was supported by the Estonian Research Council under grant PRG5614 (AIBiodivRisk) and the French National Research Agency (ANR) under grant ANR-25-CE34-0010 (BRIGlobal). Sentinel-2 imagery processed via Google Earth Engine (free academic access). CNN training labels: ESA WorldCover 2021 (ESA; CC BY 4.0) + manually corrected 8,470 training polygons (author-generated; CC BY 4.0). Hansen GFC (University of Maryland; open access). IUCN PAME database (IUCN; academic access). TVI species data from DBA paper 284 (open data; CC BY 4.0). RRPI from DBA paper 280 (open data; CC BY 4.0). Pipeline hosted on Google Cloud Platform (academic research credits). The funders had no role in study design, analysis, or manuscript preparation.

Conflict of Interest

The authors declare no conflicts of interest.

Data Availability Statement

Complete BRI pipeline source code (Python; PyTorch; FastAPI; Apache Airflow; PostGIS; MIT licence), trained CNN model weights (ResNet-50; land cover classification + deforestation detection; CC BY 4.0), October 2024 global BRI map (1 km² resolution; GeoTIFF; CC BY 4.0), all five component layer rasters (habitat change; TVI; cascade; deforestation alert; RRPI), validation datasets (247 IUCN PAME assessments; 84 BIOCHANGE sites; BRI scores), and REST API documentation are deposited in the Zenodo repository under DOI: 10.5281/zenodo.19497150. All data released under CC BY 4.0; code under MIT licence. Live BRI API (free access; academic use): <https://bri-api.biodiversity.eu> (demo endpoint; rate-limited).

Ethical Approval

This study involved exclusively computational analysis and AI model development using publicly available satellite imagery and biodiversity databases. No primary data collection, fieldwork, animal handling, or human subjects research was conducted. No institutional ethics approval was required. The BRI system is designed for conservation planning and enforcement support only; all personal data protections apply if the system is used in surveillance contexts (ranger deployment applications should be reviewed under national data protection laws).

Appendix A

BRI Pipeline Implementation and CNN Training Protocol

The Google Earth Engine script for Sentinel-2 preprocessing and the CNN fine-tuning protocol for the habitat classification and deforestation detection components are provided below.

SENTINEL-2 PREPROCESSING (GOOGLE EARTH ENGINE)

Cloud-free monthly composite: `var s2 = ee.ImageCollection('COPERNICUS/S2_SR_HARMONIZED').filterDate(start_date, end_date).filterBounds(region).filter(ee.Filter.lt('CLOUDY_PIXEL_PERCENTAGE', 20)).map(function(image) { var scl = image.select('SCL'); var mask = scl.neq(3).and(scl.neq(8)).and(scl.neq(9)).and(scl.neq(10)); return image.updateMask(mask).select(['B2','B3','B4','B8','B11','B12']); });` var composite = s2.median(); The SCL mask removes cloud (8; 9; 10) and cloud shadow (3) pixels. Select bands: B2 (blue 490nm); B3 (green 560nm); B4 (red 665nm); B8 (NIR 842nm); B11 (SWIR1 1610nm); B12 (SWIR2 2190nm) -- the 6-band combination optimised for vegetation and soil discrimination.

NDVI and change indices: `var ndvi = composite.normalizedDifference(['B8','B4']).rename('NDVI');` var nbr = composite.normalizedDifference(['B8','B12']).rename('NBR'); For bi-temporal change detection: `var ndvi_change = ndvi_t2.subtract(ndvi_t1).rename('NDVI_change');` var nbr_change = `nbr_t2.subtract(nbr_t1).rename('NBR_change');` Stack for CNN input: `var input_stack = composite.addBands([ndvi, nbr, ndvi_change, nbr_change]);` This 10-band input stack (6 spectral + NDVI + NBR + 2 change indices) provides the CNN with both spectral identity and change magnitude information for deforestation detection.

Export for CNN inference:

`Export.image.toDrive({image: input_stack, description: 'S2_monthly_composite', scale: 10, region: region, maxPixels: 1e13, fileFormat: 'GeoTIFF', formatOptions: {cloudOptimized: true}});` Process in 1 km² tiles for CNN inference on GPU: tile the exported GeoTIFF into 1024 x 1024 pixel patches (10240 m x 10240 m) with 50-pixel overlap to avoid edge effects. Pass each patch through the CNN; reassemble; aggregate to 1 km² risk score.

CNN FINE-TUNING PROTOCOL

Architecture: ResNet-50 (pretrained on ImageNet via timm Python library; open source). Modify final classification head: replace original 1000-class head with a 2-layer MLP (512 -> 11 for land cover classification; or 512 -> 1 + sigmoid for binary deforestation detection). Input: 10-band image patches (B2; B3; B4; B8; B11; B12; NDVI; NBR; NDVI_change; NBR_change); standardise each band to zero mean and unit standard deviation using training set statistics.

Training data: 8,470 manually annotated training polygons (land cover classification) + 18,470 confirmed deforestation events from Hansen GFC 2024 layer (deforestation detection). Split: 70% training; 15% validation; 15% test; spatial block split (minimum 50 km between train and test blocks to prevent spatial autocorrelation inflation of accuracy). Augmentation: random horizontal/vertical flip; random rotation (0-360 degrees); random brightness/contrast (+/-20%); random crop and resize (scale 0.8-1.2).

Training: Adam optimiser (lr = 1e-4; weight_decay = 1e-4); cosine annealing learning rate schedule (T_max = 50 epochs); batch size = 32; train for 100 epochs maximum; early stopping if validation loss does not improve for 10 epochs. Loss: cross-entropy (land cover classification); binary cross-entropy (deforestation detection; class-weighted: positive class weight = 5 to compensate for class imbalance -- deforestation events are rare relative to stable forest). Train on single NVIDIA A100 GPU; approximately 8-12 hours per training run.

Inference and BRI integration: Apply the trained CNN to monthly Sentinel-2 composites globally (processing in GEE-exported tiles; parallel inference on Google Cloud GPU instances; approximately 72 hours per monthly global run). Deforestation alert: binary classification per 10 m pixel; aggregate to 1 km² by computing proportion of deforestation-alert pixels per tile. Habitat change risk score: 1 - proportion of high-value habitat (forest; grassland; wetland; mangrove) pixels remaining in current vs. reference (baseline 2015-2020 mean composite) month. Feed both scores into the BRI composite computation (weights from Table 1).