2026 International Mathematical Modeling Challenge (IM^{2}C)

§ Summary

Etosha National Park is too vast, geographically diverse, and limited in operational resources to be protected uniformly. Due to the extremely dry conditions and limited water supply, wildlife is unevenly dispersed, tending to gather around waterholes. Threats vary as well, with poaching risks increasing around intrusion points and wildfire danger particularly concentrated on the western half (Turner et al., 2022). On top of that, detecting a threat is only meaningful if rangers can actually respond quickly enough to make a difference. In this context, we define *protection* as **the fraction of expected ecological loss that is prevented**. This turns the task into an optimization problem: given limited personnel, patrol hours, drones, and sensors, where should resources be placed, and how can we allocate them, to minimize harm?

To answer that question, we built a geographic optimization model grounded in four main components:

- a **park-wide spatial layer** encompassing roads, camps, waterholes, terrain, and fire hazards. This data was extracted from open source databases like NASA FIRMS and POWER, giving the rest of the model a foundation to build on.
- **ecological consequence and threat heatmaps** that identify the areas of highest vulnerability and ecological value. After normalizing environmental and operational variables, these layers were combined into a baseline loss surface showing regions of the park facing the greatest expected ecological harm.
- a **terrain-aware mobility model** that integrates slope-adjusted road travel and Tobler's hiking function for off-road traverse to estimate how quickly each section of the park can be reached. Travel time is then translated to response success through a logistic curve.
- a **mixed-integer optimization model** that creates a deployment plan based on a package of patrol routes, drone sorties, and sensor sites. Using a bounded concave piecewise-linear surrogate for detection, the probabilistic overlap of patrols, drones, and sensors was incorporated within a Mixed Integer Linear Program (MILP) to maximize avoided loss under staffing constraints.

Under the baseline package of **10 patrol routes, 4 drone sorties, and 12 sensor sites**, the model recommends a hotspot-first deployment strategy concentrated near waterholes and road-reachable corridors. The estimated protection indices (fraction of loss avoided) are 0.148 for dry-season poaching, 0.136 for wet-season poaching, and 0.129 for the combined all-threat scenario. However, within the highest-risk 10% of grid cells, the protection index is substantially higher at 0.250, showing that the model prioritizes depth over breadth. The operational plan requires 83 active personnel, which converts to about 291 total staff, placing it just below the park's 295-person ceiling. The staffing curve shows diminishing returns, with protection beginning to plateau near 167 active personnel, equivalent to 595 total staff.

Robustness tests show that under fixed baseline deployment, protection in the combined case gradually levels off. This indicates that the model should be used as a strategic planning tool that is rerun as seasonality, staffing, and threats change. Lastly, the same mathematical structure can be transferred to other parks, such as Yellowstone and Kaziranga, as long as the ecological features, mobility assumptions, and operational constraints are recalibrated locally.

§ Letter to IMMC

Dear IMMC,

Five days ago, you handed us a challenge worthy of a very sleep-deprived team of three student mathematicians: Build a model to help guide protection strategies for large wildlife reserves, starting with Etosha National Park. We're glad to report that in this remarkably short span of time, our team was able to develop a robust mathematical model for limited resource allocation and turn it into a complete and actionable deployment strategy for Etosha National Park. What sets Etosha apart from other reserves is its scarce water supply. As such, the park's waterholes serve as lifelines for wildlife, pulling animals into concentrated, high-stakes zones. That is why our strategy deliberately deprioritizes some areas to deploy patrols, drones, and sensors in the most vulnerable zones where ecological importance, danger, and rapid response overlap.

To build out this plan, we divided Etosha into 996 grid cells, each measuring 5 km by 5 km, and evaluated every cell by asking four practical questions: Which cells are the most ecologically valuable? Which areas are most exposed to threats? Where are incidents most likely to be detected? And where can response teams arrive in time? Our model considers a wide range of factors –including terrain, proximity to gates and fences, wildfire-prone areas, and roads – to determine the inherent value of each cell. Crucially, a threat that is seen but cannot be reached in timely manner is not actually worth protecting, so response speed is ingrained in the model from the very start.

As mentioned in our summary, our deployment plan relies on a baseline package of 10 patrol routes, 4 drone sorties, and 12 sensor sites, which is a realistic workload for Etosha's 295 staff members. Under the combined scenario, our hotspot-prioritizing plan is estimated to prevent about 13% of expected ecological harm across the park and 25% of expected harm within the most vulnerable 10% of locations. While these figures may not seem staggering at first glance, they are meaningful precisely because the park is so vast and resources are so limited. In fact, that margin of protection can mean the difference between survival and local extinction for endangered species such as the black rhino.

To put the effect of staffing into perspective, our model shows that increasing Etosha's workforce to 417 total staff would raise the park-wide loss-prevention rate to approximately 16%–a meaningful improvement, but one that requires a 43% increase in personnel. Beyond 585 staff, however, our model shows clear diminishing returns, where additional rangers are largely sent to areas that are low-risk or so remote that slow response times make intervention unlikely anyway. Lastly, our model found that staffing shortfalls were the single most influential factor affecting protection, with a 30% reduction in staff corresponding to a 15.8% decline in protection.

Overall, we recommend that Etosha adopt the hotspot-first baseline strategy proposed, prioritize staff and rapid response, and rerun the model whenever seasons, staffing levels, road access, or threat patterns change. Because the program also reran successfully on Yellowstone and Kaziranga after local recalibration, we believe our model is a credible and transferable method for parks in general.

Attached below is an infographic of our Etosha deployment plan and the benefits it provides. We would like to extend our gratitude to IMMC for giving us a chance to showcase our mathematical modeling, critical thinking, and problem solving skills to tackle real-world challenges.

Yours sincerely,
Team 2026025

TEAM2026025
ETOSHA PROTECTION PLAN OVERVIEW

Recommended operating idea: Focus protection near the waterhole-linked and road-reachable hotspot corridors, where wildlife concentration, illegal access, and realistic response opportunity overlap most strongly.

Baseline Deployment Map

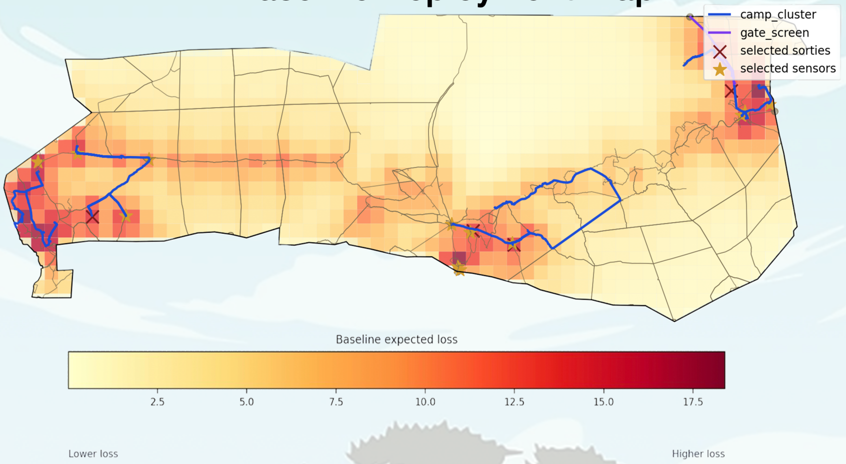

Figure 1: Patrols, drone anchors, and sensor sites cluster around the western, central, and eastern hotspot corridors rather than being spread uniformly across the full park.

Deployment Package

10 Patrol Routes

4 Drone Sorties

12 Sensor Sites

12 Base Staff

Active staff total
10x6+4x2+12x0.25+12 = **83**

Protection Outcome (PI)

Combined 0.129

Dry Season 0.148

Wet Season 0.136

Hotspot Protection* 0.250

The strategy protects hotspot corridors much more strongly than the park as a whole.

Management Signal

83 Active Personnel
= 291 total staff corresponds to the current protection outcomes.
This is the best use of the 295 limited staff.

119 Active Personnel
= 417 staff reaches the strong planning threshold of combined protection benchmark **0.16**

167 Active Personnel
= 585 staff is where the current model begins to plateau near its best protection **0.18**

One Minute Message

Etosha should not be patrolled evenly. The strongest current strategy is to concentrate patrols, drones, and sensors where wildlife value is highest, threat pressure is strongest, and response is still realistically fast. With only 295 staff, the park cannot reach an ideal protection level everywhere (585 staff). That means the goal is not perfect coverage, but limiting the most serious harm. The key is to place resources deliberately and preserve fast response where protection matters most.

§ Table of Contents

§ Introduction

National parks represent one of humanity's biggest pledges to protecting nature. Having been established in nearly 100 countries around the world, these protected zones play a major role in saving endangered ecosystems and preserving biodiversity over long periods of time ("Protected Areas," n.d.).

With each passing day, however, the vision of national parks is becoming more and more elusive. Environmental research has shown that ecological protection in parks can be incomplete, unevenly enforced, and at times more reactive than proactive (Pulido-Chadid et al., 2023). Reserves such as Etosha National Park contain vast and complex terrains, but unfortunately, conservation resources remain limited while threats like poaching and wildfires are widespread. Emerging technologies, including drones, sensors, and satellite data, can be powerful tools for protection, but due to their scarcity, their effectiveness depends heavily on their strategic deployment (van Doormaal et al., 2021).

To address this issue, this paper will use a mathematical modeling approach to optimize resource allocation in national parks, beginning with Etosha as a case study. The central question we will set out to answer is: how should a park allocate limited patrols, drones, sensors, and staff to reduce expected ecological harm as much as possible?

§1. Problem Framing

The key consideration in the 2026 IMMC problem is developing a quantitative notion of protection that is realistic and effective for a large reserve like Etosha with limited personnel, diverse threats, and substantial **spatial heterogeneity** (Ghoddousi et al., 2022). Crucially, a model based on uniform geometric coverage alone fails to address this concern. The problem with prioritizing coverage is that it treats all space as equally important and ignores the possibility of seeing a threat and still being unable to reach it in time. Those simplifications would be particularly harmful in Etosha, where ecological importance is concentrated around waterholes and operational success is linked to the road network and terrain circumstances.

We therefore present a **probabilistic model**. First, we define protection as the *reduction of expected ecological loss*. In order to quantify this at specific points within the park, we partition the map into 5 km $\times$ 5 km square grid cells (996 in total), with g being the index of a grid cell and A_g being the cell area. Setting C_g as ecological consequence, λ_g as threat intensity, R_g as response success, and D_g as detection success (see Appendix C for the detailed notation list), we define the baseline loss (loss that would be incurred if there were no intervention) in cell g to be

$$L_g^0 = A_g\, C_g\, \lambda_g,$$

and the residual loss after deployment to be

$$L_g = A_g\, C_g\, \lambda_g\, (1 - D_g R_g) = L_g^0(1 - D_g R_g).$$

Aggregating over the park yields the total baseline loss and total post-deployment loss, respectively:

$$L_0 = \sum_{g\in G} L_g^0, \qquad L = \sum_{g\in G} L_g.$$

Finally, we define the **protection index** to be

$$PI = 1 - \frac{L}{L_0}.$$

Central Idea — The ultimate objective is to minimize

$$E[\text{ecological loss}],$$

which is equivalent to maximizing the protection index PI for a given scenario.

To do this, each cell g is evaluated by combining four factors: its ecological importance (C_g); its vulnerability to poaching, intrusion, and fire (λ_g); the probability of detecting harm in the cell(D_g); and the probability that the protection system can intercept a threat(R_g).

Note: in the MILP that we use (to be introduced later), the detection probability D_g will be replaced by a bounded piecewise-linear surrogate S_g. Until Section 6, we will write D_g to denote the conceptual detection probability and then use S_g to refer to the surrogate.

This particular modeling approach creates three important mathematical advantages:

1. It prioritizes certain regions of Etosha over others; this hotspot resource allocation approach is essential due to scarcity
2. It distinguishes harm observation from actual intervention and interception, allowing us to model detection and response separately
3. It supports the scenario and sensitivity analysis that we will later conduct since response and detection are probabilistic terms

As alluded to in the summary, the main components of the model include: a waterhole-centered ecological value layer, a hazard layer that combines satellite fire context with poaching data, and a hybrid grid-graph layout that distinguishes continuous biological risk from terrain-aware road mobility.

§2. Etosha Context and Data Resources

Etosha National Park spans roughly 22,935 square kilometers (Government of Namibia, n.d.). Due to this sheer size, our model uses the park's actual layout of waterholes, roads, gates, and camps to determine where protection should be concentrated. We obtained real world data: park boundaries and road networks from OpenStreetMap (OSM) (OpenStreetMap contributors, n.d.; Haklay & Weber, 2008); 56 verified waterholes, 6 camps, and 4 gates (*Map of Etosha,* n.d.); a 30m-class elevation raster from the OpenTopography global DEM service (`COP30`) (OpenTopography, 2021; Copernicus Programme, n.d.); and wildfire context from NASA FIRMS (NASA, n.d.-a) and NASA POWER (NASA, n.d.-b; NASA, 2025).

Note: The full data, along with the foundational geometric maps used to build the value and threat fields, is deferred to Appendix A.

§ Mathematical Model

Note: The majority of key mathematical notation will be introduced inline. Nonetheless, a comprehensive list of variables can be found in Appendix C.

§1. Hybrid Grid–Graph Representation of the Park

First, the optimization model requires us to break down the vast landscape of Etosha into smaller areas. Due to the presence of both continuous terrain and discrete roads, we choose to employ a **hybrid grid–graph representation**: the grid is used to represent how ecological value and threats are distributed across individual cells, while the graph is used to represent roads and shape response design.

The **grid** is the model's primary component. We will subdivide the park into square cells of size 5 km $\times$ 5 km each and index them with numbers $g \in G$. Refer to Subsection 1. (Problem Framing) for more quantities stored in each grid cell. This grid-based representation is effective because ecological value and threat intensity are not confined to roads or points.

The **graph** serves a different purpose, measuring how reachable each part of the park is from available response bases. Camps, gates, and road junctions form the nodes of a terrain-aware road graph, and edges are weighted by travel time rather than geometric distance. This allows the model to compute best-response times and convert them into response success probabilities. In other words, the graph captures operational mobility, rather than ecological exposure.

For each cell, we compute its area, centroid, and distances to the nearest waterhole, road, camp, gate, and boundary. These quantities feed directly into the value, threat, coverage, and response frameworks we develop below.

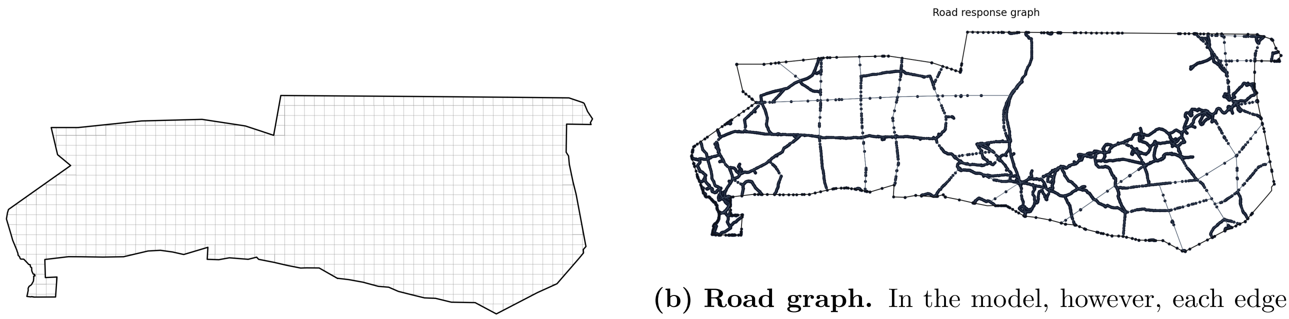

(a) Grid layout. 996 cells, each 5 km $\times$ 5 km.

(b) Road graph. In the model, however, each edge is weighted by its estimated travel time instead of its geometric length.

Figure 1: Hybrid spatial structure. The grid stores ecological value and threat, while the road graph dictates movement and response.

§2. Ecological Value Layer

The ecological value layer answers the question: *Which areas are most costly to lose?* In Etosha, this is fundamentally impacted by the distribution of water sources. Dry-season wildlife is concentrated near permanent water, whereas in the wet season, wildlife migrates outward as short-term water resources become accessible.

> **Assumption** (1)**.** We assume that it is not feasible to track individual animals or simulate their movement paths. Instead, we represent ecological value as a season-specific spatial density field. This is a necessary methodological choice because high-resolution telemetry data is unavailable for the entire park population, and simulating multi-agent movement is computationally intractable at the strategic allocation scale. Ultimately, what matters most for long-term planning is broad patterns of concentration around habitat features like waterholes.

Normalized to $[0, 1]$, consequence surfaces are built from season-specific kernels that weight species accordingly. Appendix D2 gives the full details and derivations.

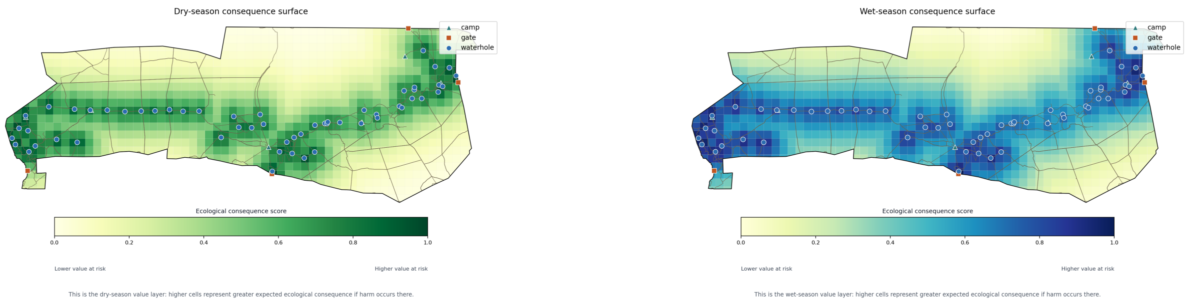

(a) Dry-season consequence

(b) Wet-season consequence

Figure 2: Consequence surfaces: dry-season value collapses toward the waterhole network while Wet-season value spreads more broadly across the park.

§3. Threat Layer

The threat layer answers the complementary question: *Where is harm most likely to occur?* This is an important consideration because ecological value alone cannot fully determine the need for intervention; the threat must also be quantified explicitly. In Etosha, we model each of the three main threats (poaching, intrusion, and fire) separately since they originate differently. Poaching is driven by ecological value and dependent on human access; intrusion depends primarily on human access; and, fire is caused by a combination of weather conditions and terrain.

For each type of threat, we model risk using a relative spatial intensity surface rather than forecasts of distinct incident counts within each cell. The feasible model is designed to identify areas in which poaching, intrusion, or fire is relatively more likely, not to discretely predict the inherently uncertain number of future incidents.

> **Assumption** (2)**.** We assume that offenders (poachers, intruders, arsonists, etc.) are boundedly rational. We presume that they are aware of ecological value distribution, the locations of roads and gates, and the difficulty of the terrain; however, they do not know the real-time deployment of protection resources (rangers, drones, sensors, etc). This is reasonable because these geographical features are publicly available, whereas live patrol and drone locations are not generally observable. This assumption helps avoid additional game-theoretic complications.

Threat surfaces are built as normalized weighted sums of bounded spatial kernels. For poaching, risk is the highest where ecological value and human access are both high; for intrusion, risk scales with human access; for fire, structural exposure, satellite detections, weather, access pressure, and terrain are combined to determine risk.

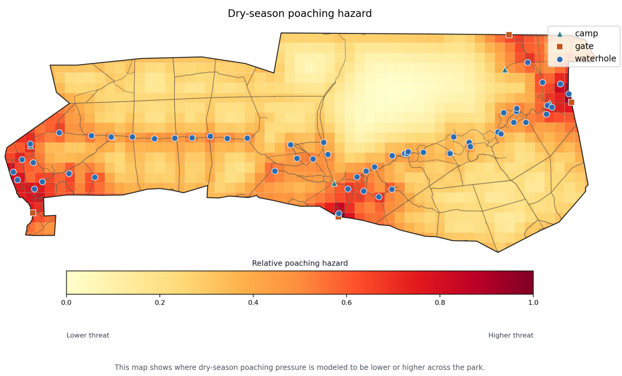

(a) Dry-season poaching hazard

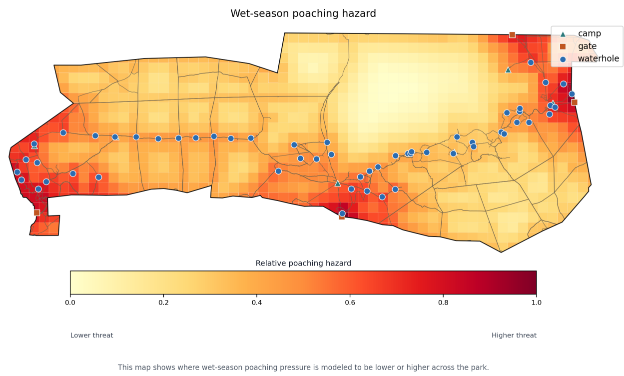

(b) Wet-season poaching hazard

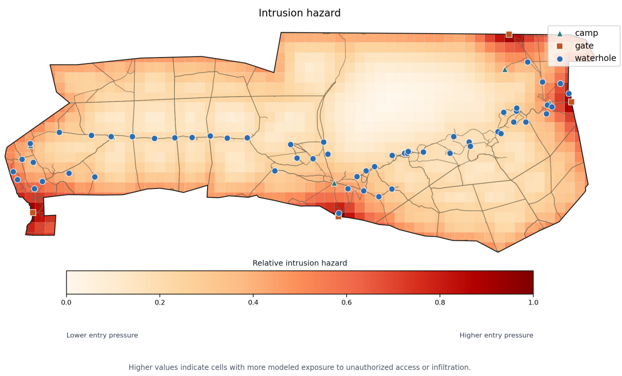

(c) Intrusion hazard

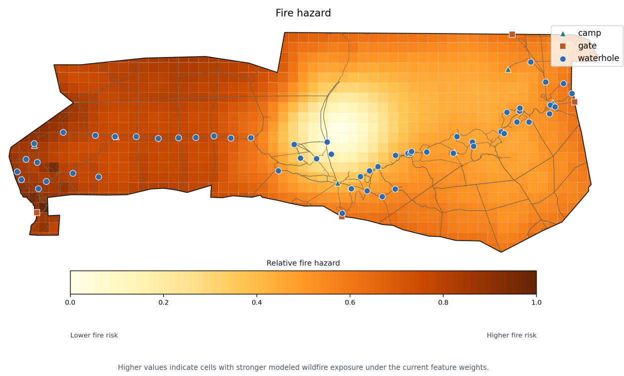

(d) Fire hazard

Figure 3: Threat fields expressed on the same normalized hazard scale.

The ecological value layer identifies areas where ecological damage would be most severe if an incident occurred, while the threat layer identifies areas where incidents are most likely to occur. Neither is sufficient to quantify loss on its own: high-value areas may be relatively secure, and conversely, high-threat areas may be less ecologically significant. Our model defines baseline loss in cell g as

$$L_g^0 = A_g C_g \lambda_g,$$

where ecological value and threat intensity are both taken into account.

§4. On-road and Off-road Mobility

Protection response in Etosha depends not only on distance, but also on terrain, road slope, and other factors that influence how quickly teams can actually move through the park. Therefore, we model mobility using terrain-adjusted travel time in addition to geometric distance.

Let θ_{ij} denote the local slope angle along road segment (i, j), and let $|\tan\theta_{ij}|$ be the corresponding steepness. On-road travel speed is modeled as

$$v_{ij}^{\text{road}} = v_{\text{class}(ij)}\,\text{clip}(1 - c|\tan\theta_{ij}|,\ f_{\min},\ 1)\,,$$

where:

- $v_{\text{class}(ij)}$ is the base speed for the road class of segment (i, j),
- c is a slope-penalty coefficient,
- $f_{\min}$ is a minimum speed factor that prevents moderate slopes from driving speeds to unrealistically small values.

This form is deliberately simple. It preserves the key operational fact that steeper roads are slower to traverse, while still keeping road travel tied to the underlying road hierarchy. In other words, a main road remains faster than a connector road, but both are penalized when slope increases.

> **Assumption** (3)**.** We assume that large-scale response performance is governed primarily by base locations, road connectivity, and terrain-adjusted travel time. This is appropriate for strategic planning because these are the main determinants of whether a team can reach a threat in time, whereas short-term disruptions such as vehicle breakdowns, flooding, or temporary blockages are more difficult to calibrate reliably. The assumption lets the model translate geography directly into response quality.

Once defenders leave the road network, movement is modeled with *Tobler's hiking function*:

$$v^{\text{off}}(\theta) = 6\exp\Big(-3.5|\tan(\theta) + 0.05|\Big),$$

where θ is the local off-road slope angle and $v^{\text{off}}(\theta)$ is measured in km/h.

Tobler's law is suitable here because it is a standard terrain-based walking-speed model. It captures the basic fact that off-road travel is fastest on gentle downhill slopes and slows on both steep uphill and steep downhill terrain. Since the final approach to a target cell often cannot remain entirely on the road network, this provides a reasonable approximation for terminal on-foot travel.

If base b reaches cell g by road travel followed by a final off-road segment, then total response time is

$$\tau_{bg} = \tau_{bg}^{\text{road}} + \tau_{bg}^{\text{off}},$$

where:

- τ_{bg}^{road} is shortest-path road travel time from base b to the best reachable road access point for cell g,
- τ_{bg}^{off} is the final off-road travel time from that access point to the cell.

The best-response time for cell g is then $\tau_g = \min_b \tau_{bg}$. Thus τ_g is the fastest attainable response time over all available bases.

We convert this travel-time quantity into a response-success probability through the logistic function

$$R_g = \frac{1}{1 + \exp\Big(\kappa(\tau_g - \tau_0)\Big)},$$

where:

- τ_0 is the midpoint response time,
- κ controls how sharply response quality falls as travel time increases.

In the present calibration (Appendix K), $\tau_0 = 2.000\,\text{h}$ and $\kappa = 2.2$, so response probability declines quickly once the fastest available response exceeds roughly two hours.

The logistic form is useful because both parameters have direct operational meaning. First,

$$R(\tau_0) = \frac{1}{2},$$

so τ_0 is the point at which interception changes from more likely than not to less likely than not. Second,

$$\frac{dR}{d\tau} = -\kappa R(\tau)\Big(1 - R(\tau)\Big),$$

which implies that the curve is steepest at $\tau = \tau_0$, where

$$\left.\frac{dR}{d\tau}\right|_{\tau_0} = -\frac{\kappa}{4}.$$

This is exactly the behavior we want: small improvements in travel time matter most near the operational breakpoint, where a cell is close to being either realistically reachable or too late to defend.

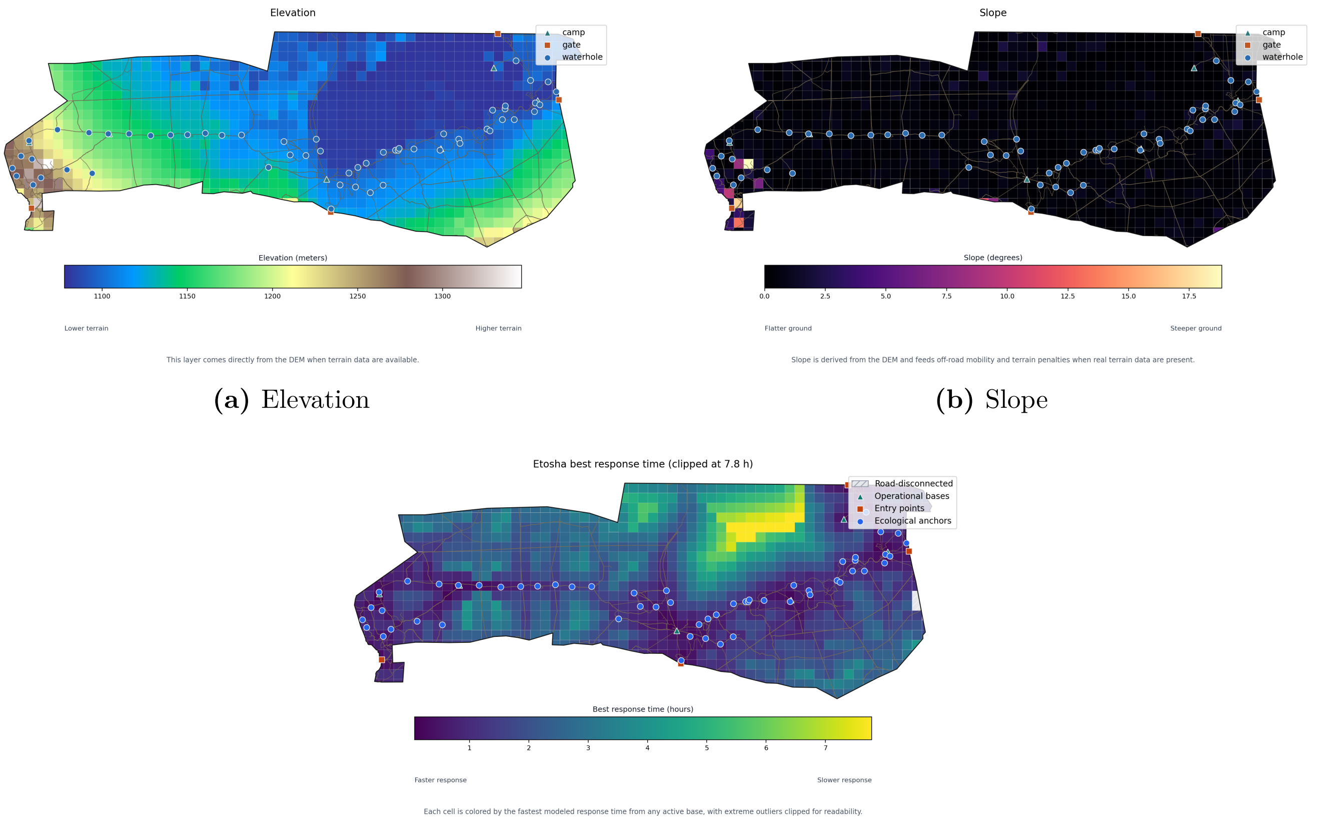

(a) Elevation

(b) Slope

(c) Best response time

Figure 4: Terrain-aware mobility. Elevation and slope are converted directly into travel-time differences, which then determine response quality across the park.

The computed response-time map is informative as it shows which parts of the park can be reached quickly and which are harder to access. The median best-response time is 1.940 h; 89.600% of cells lie within 4.000 h; 96.600% lie within 6.000 h; and only 3 cells appear effectively unreachable under the current road-network+off-road-trek model. These outliers help explain the fairness gap discussed later.

§5. Patrol, Drone, and Sensor Coverage

Patrols, drones, and sensors each serve distinct functions. The model therefore does not treat them as equivalent assets; instead, each asset type is assigned an influence matrix.

> **Assumption** (4)**.** We assume that each asset type is homogeneous: that is, the performance of all patrols are equal, all drones are equal, and all sensors are equal. This is a fair assumption since for this modeling task, overall deployment structure matters more than small performance fluctuations. It will come in handy by helping keep the MILP interpretable by allowing each asset class to be modeled with common parameters.

§5.1 Patrol routes

Patrols are drawn from a route library consisting of roads, camps and gates. The current Etosha library includes:

- camp-cluster patrols,
- gate-screen patrols, and
- perimeter-screen patrols.

If Γ_r and c_g denote polyline route r and the centroid of cell g, respectively, then the implemented patrol influence is

$$\phi_{rg} = \alpha_P \exp\left(-\frac{\mathrm{dist}(c_g, \Gamma_r)}{\rho_P}\right) \cdot \mathbf{1}\left\{\alpha_P \exp\left(-\frac{\mathrm{dist}(c_g, \Gamma_r)}{\rho_P}\right) \geq \varepsilon_P\right\}.$$

Operationally, this says a patrol route affects not only the cells it passes through, but also nearby cells, with influence decaying exponentially as the distance to the centroid increases. The threshold ε_P truncates very small influence values, which keeps the matrix sparse and preserves computational efficiency.

§5.2 Drone sorties

A drone sortie is a planned drone mission, basically one scheduled flight launched around a selected anchor point to scan a surrounding area for a limited time. Drone sorties are anchored near high-risk regions and contribute broader but time-limited influence. If sortie d is anchored at point a_d and has coverage radius R_d, then

$$\psi_{dg} = \alpha_D \max\left\{0, 1 - \frac{\|c_g - a_d\|}{R_d}\right\} \cdot \mathbf{1}\left\{\alpha_D \max\left(0, 1 - \frac{\|c_g - a_d\|}{R_d}\right) \geq \varepsilon_D\right\}.$$

This linear taper gives drones a wider but shallower footprint than patrols, which matches their role as flexible reach assets rather than persistent corridor coverage.

§5.3 Sensor sites

A sensor site, contrary to drones, is a fixed monitoring position. Sensors are sited at camps, gates, selected waterholes, and strategic interior points. Their influence is based on line of sight, distance decay, and a finite operating radius. If $\mathrm{vis}_{sg} \in \{0, 1\}$ records whether cell g is visible from site s and R_S is the maximum sensor radius, then the implemented sensor influence is

$$\chi_{sg} = \mathrm{vis}_{sg} \max\left\{\chi_{\min}, \chi_0 \exp\left(-\frac{\mathrm{dist}(s, c_g)}{\rho_S}\right)\right\} \mathbf{1}\{\mathrm{dist}(s, c_g) \leq R_S\}.$$

Because line of sight is terrain-dependent, the present analysis uses the DEM to build the binary visibility term vis_{sg} before the distance-decay term is applied.

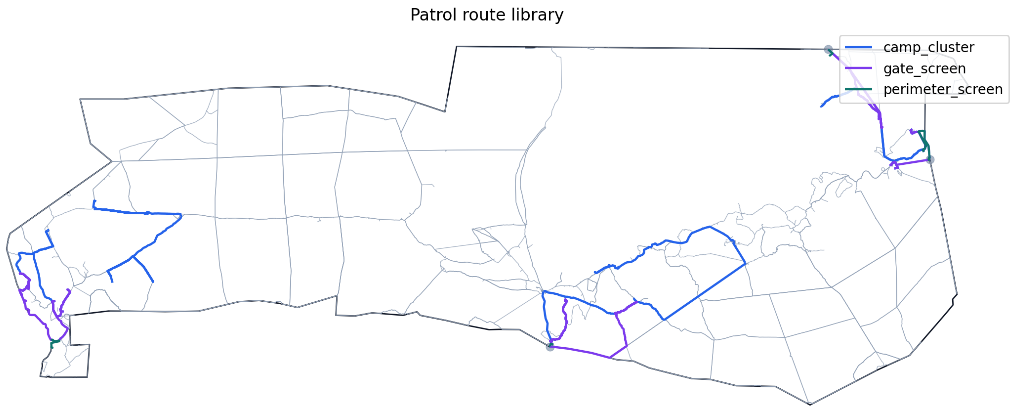

(a) Patrol route library

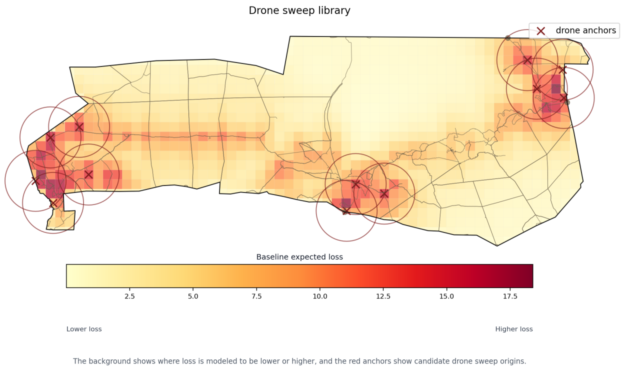

(b) Drone sweep library

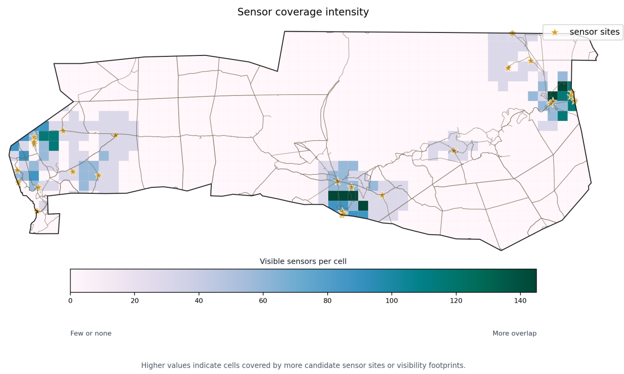

(c) Sensor coverage intensity

Figure 5: Patrol, drone, and sensor coverage libraries. The three assets differ in movement, reach, and influence, so they must be considered separately.

The following matrices link the park's spatial geometry to the optimization model:

$$\phi_{rg} \quad \text{(patrol influence)}, \qquad \psi_{dg} \quad \text{(drone influence)}, \qquad \chi_{sg} \quad \text{(sensor influence)}.$$

Each matrix measures how strongly a given asset affects a given cell. The matrices are built by assigning each asset-cell pair an influence coefficient according to the operational capacity of that asset. Specifically, patrol influence depends on distance from the patrol route, drone influence is determined by aerial coverage range, and sensor influence is derived from the line of sight and distance. Together, the three assets cover a wide range of protection geometries.

§6. Detection and Interception Model

This section introduces the main simplification that makes the model solvable. Up to this point, detection has been written conceptually as D_g, the probability that harmful activity in cell g is detected. If patrol, drone, and sensor detections are conditionally independent at the cell level, then the exact detection probability is

$$D_g^{\text{exact}} = 1 - \left(1 - P_g^{\text{patrol}}\right)\left(1 - P_g^{\text{drone}}\right)\left(1 - P_g^{\text{sensor}}\right).$$

This is the natural conceptual formula, but it is nonlinear in the asset-selection decisions, so optimizing it directly would lead to a mixed-integer nonlinear program rather than a MILP.

Expanding the product gives

$$\begin{aligned} D_g^{\text{exact}} = P_g^{\text{patrol}} &+ P_g^{\text{drone}} + P_g^{\text{sensor}} \\ &- P_g^{\text{patrol}} P_g^{\text{drone}} - P_g^{\text{patrol}} P_g^{\text{sensor}} - P_g^{\text{drone}} P_g^{\text{sensor}} \\ &+ P_g^{\text{patrol}} P_g^{\text{drone}} P_g^{\text{sensor}}. \end{aligned}$$

The first line is additive, while the remaining terms capture overlap and saturation. Those higher-order terms are what make the exact model extremely difficult to optimize.

To preserve linearity, we aggregate all selected monitoring effort in cell g into the scalar

$$M_g = \sum_r \phi_{rg} y_r + \sum_d \psi_{dg} z_d + \sum_s \chi_{sg} x_s,$$

where $y_r, z_d, x_s \in \{0, 1\}$ indicate whether patrol route r, drone sortie d, or sensor site s is selected. The quantity M_g is therefore the total monitoring intensity in cell g.

We then map M_g into a bounded, concave, piecewise-linear surrogate S_g, which is the implemented MILP proxy for detection (see Appendix D). Thus, D_g remains the conceptual detection probability, while S_g is the quantity actually used in the optimization model.

This replacement is reasonable because, for small to moderate component probabilities,

$$D_g^{\text{exact}} = P_g^{\text{patrol}} + P_g^{\text{drone}} + P_g^{\text{sensor}} \; + O\Big(P_g^{\text{patrol}} P_g^{\text{drone}} + P_g^{\text{patrol}} P_g^{\text{sensor}} + P_g^{\text{drone}} P_g^{\text{sensor}}\Big),$$

so the additive part is first-order, while overlap enters only at higher order. The surrogate S_g preserves the desired behavior that detection rises with monitoring but with diminishing marginal gains. Similar to previously mentioned (just with S_g instead of D_g this time), the resulting residual loss for each cell is

$$L_g = L_g^0 - L_g^0 R_g S_g.$$

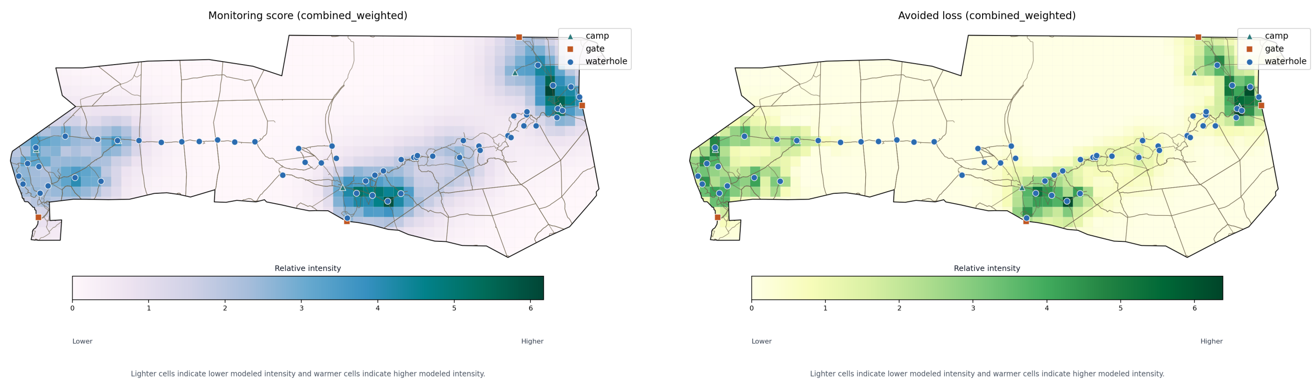

(a) Monitoring score M_g **(b)** Avoided loss

Figure 6: Monitoring itself is not equivalent to protection. The model values monitoring only where it can be converted into avoided loss through timely response.

With baseline loss coefficients, response probabilities, and coverage matrices precomputed, the strategic allocation stage becomes a MILP. If u_ℓ denotes the ℓth supporting line of the concave detection envelope for S_g, then the implemented Etosha baseline solve is summarized compactly by the following statement.

Proposition (Allocation statement)
After preprocessing, the strategic allocation problem is a mixed-integer linear program. In the Etosha baseline, the hotspot-monitoring floor is active while the remote-sector fairness floor is left inactive.

Core strategic allocation problem

$$
\begin{aligned}
\max_{x,y,z,M,S} \quad & \sum_{g \in G} L_g^0 R_g S_g \\
\text{s.t.} \quad & M_g = \sum_r \phi_{rg} y_r + \sum_d \psi_{dg} z_d + \sum_s \chi_{sg} x_s && \forall g \in G, \\
& S_g \le u_\ell(M_g) && \forall g \in G,\ \forall \ell = 0, \dots, L-1, \\
& \sum_r n_r y_r \le B_{\text{patrol}}, \qquad \sum_r T_r y_r \le H_{\text{patrol}}, \\
& \sum_d m_d z_d \le B_{\text{drone}}, \qquad \sum_d T_d z_d \le H_{\text{drone}}, \\
& \sum_s x_s \le B_{\text{sensor}}, \\
& \sum_{g \in H} M_g \ge \eta_H |H|, \\
& x_s, y_r, z_d \in \{0,1\}, \qquad M_g \ge 0, \qquad 0 \le S_g \le v_L.
\end{aligned}
$$

Appendix D1 gives the full derivation from the nonlinear detection model to this tractable surrogate.

§7. From Identification to Resource Allocation

The earlier sections define protection, construct the spatial layers (value, threat, etc.), and formulate the MILP. We now use these components that we have derived to build a resource allocation model for Etosha National Park. More technical derivations are deferred to Appendix D2, while the full calibration and result ledgers are deferred to Appendix E, Appendix F, and Appendix K.

§8. Calibration and Resource Limits

Before it can allocate anything, the model needs two more kinds of inputs: firstly, a calibrated description of Etosha, and secondly, a realistic set of resource and asset limits that the MILP solver can work with.

The baseline case scenario uses 5 km × 5 km cells, the terrain-aware response surface introduced earlier, a logistic response curve with $\tau_0 = 2.0$ h and $\kappa = 2.2$, and the bounded detection surrogate described in Section 6 and formalized in Appendix D2. For the combined scenario, the threats dry poaching, wet poaching, intrusion, and fire are weighted 0.40, 0.35, 0.15, and 0.10, respectively (see Appendix K). Also, we define the *hotspots* to be the top 10% of cells with the highest baseline loss.

> **Assumption** (5)**. Baseline resource limit.** We assume a baseline deployable resource limit of 10 patrol routes, 4 drone sorties, and 12 sensor sites. Since the main focus is to decide *how to assign* patrols, sorties, and sensor sites, it is more important for this choice to be realistic. We will show later that for this choice, the scale is close to Etosha's staffing ceiling (295) and that it balances ground presence, aerial span, and persistent sensing.

> **Assumption (6). Standardized operating units.** We assume that patrols are modeled as 6-person units, drone sorties as 2-person units and sensor support as 1 staff member per 4. We also assume that 12 background staff are required to keep operations running (e.g. for dispatch, logistics, admin support throughout) throughout. This assumption will keep the allocation model focused on deployment structure rather than on small differences between vehicles, crews, or devices.

These assumptions are extremely important because otherwise, the model would need a second optimization layer for how many assets to acquire in addition to where/how to deploy them. For this paper, we instead fix a realistic resource package and optimize its use, which makes the results easier for the park to interpret since the MILP solver can now make tangible conclusions about where the resources should be deployed.

Several of the above calibration choices also follow directly from the structural features of Etosha; for example, when determining the weight of each threat. In the dry season, wildlife clusters around permanent water, so the ecological value is more concentrated compared to the wet season. Access kernels are centered on roads, gates, and fences because those are the main entry and exit points. The poaching and response modules are therefore the most reliable parts of the calibration; the fire module is more approximate and is best understood as just a simplified hazard proxy. Detailed parameters are provided in Appendix E and Appendix K.

> **Note on MILP optimization goal**
>
> The numbers 10/4/12 are assumed baseline limits. The MILP then chooses the best *specific* patrol routes, drone sorties, and sensor sites within those limits. Note that the reported "optimal PI" is the highest feasible PI under the chosen resource limits.

§9. Baseline Allocation Results

The combined baseline loss layer shows where consequence and threat overlap. Meanwhile, the response-time layer shows where those losses can still be reached quickly enough for intervention to matter substantially.

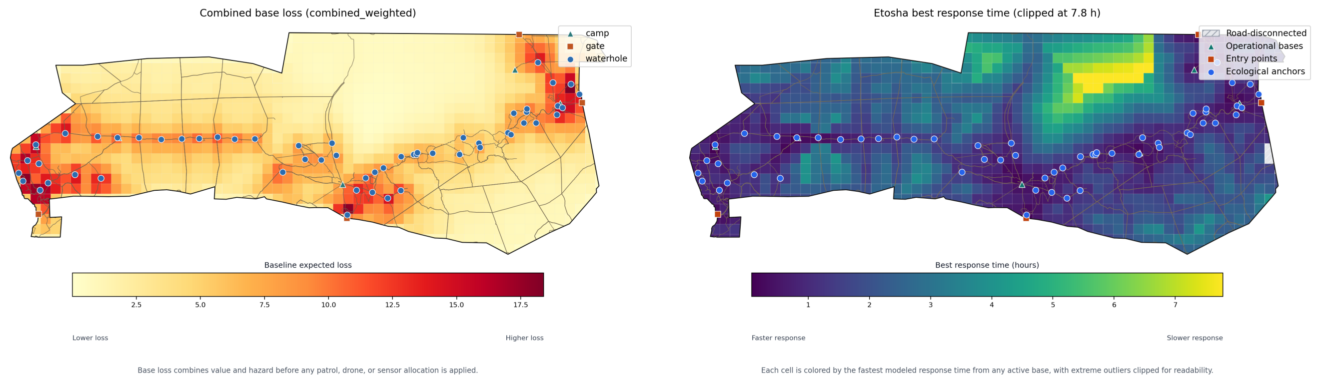

(a) Combined baseline loss

(b) Best response time

Figure 7: Loss and response layers placed in conjunction

Together, these two surfaces explain the allocation problem well; that is, instead of attempting to watch every cell equally, the model tries to place the scarce monitoring we have available where it can still be converted into avoided loss. According to the current calibrations, the median best-response time is 1.940 h; 89.600% of cells lie within 4.000 h; 96.600% lie within 6.000 h; and only 3 cells are effectively isolated. Those outliers represent the trade-off in this hotspot conservation approach.

Table 1: Scenario-level Etosha results

Scenario	Protection index	Avoided loss	Hotspot PI
Dry poaching	0.148	506.0	0.244
Wet poaching	0.136	570.4	0.253
Combined weighted	0.129	534.4	0.250

Across all three scenarios, this table reveals a pattern: park-wide protection is moderate, but hotspot protection is much higher. This discrepancy is the defining factor of our strategic allocation model under scarcity: before it spreads resources on low-return areas, the model protects the cells in which each extra unit of effort prevents the most expected loss.

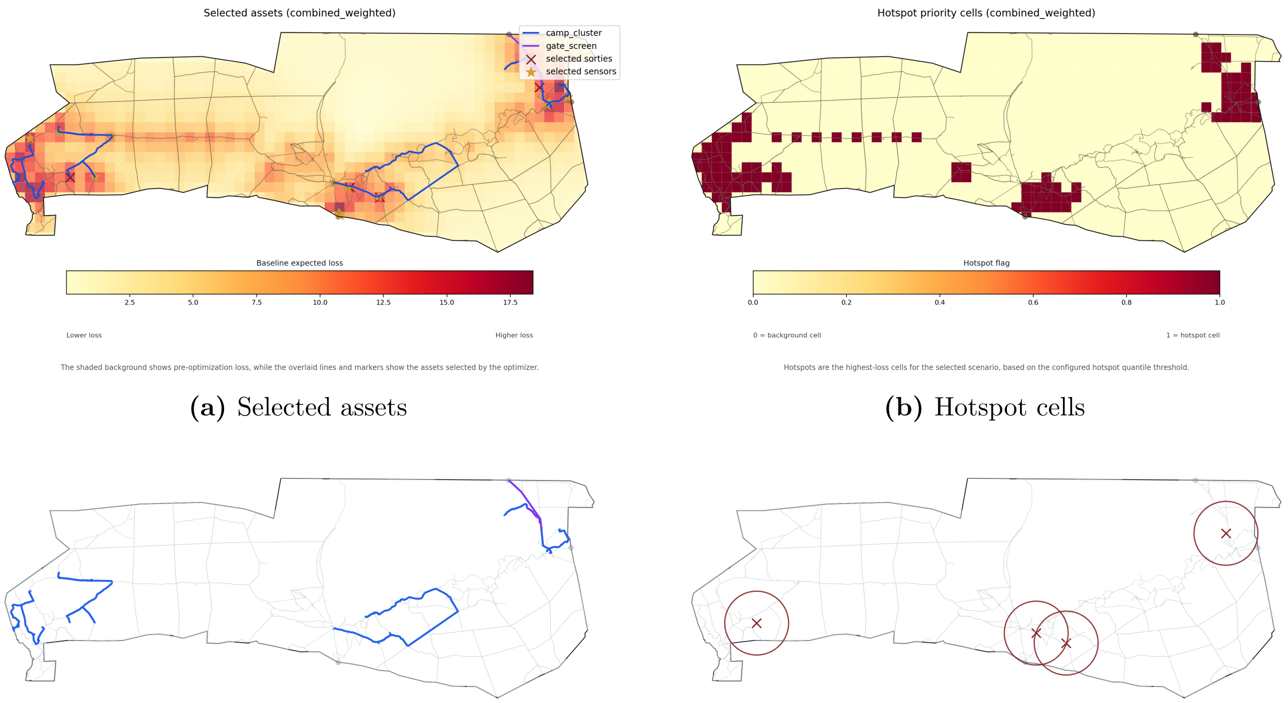

(a) Selected assets

(b) Hotspot cells

(c) Selected patrol geometry

(d) Selected sortie geometry

Figure 8: Recommended optimal deployment. The selected assets cluster around waterhole-linked and response-reachable hotspot corridors. At the same time, the other graphs show the individual patrol and sortie geometries chosen in the baseline recommendation.

The deployment pattern that we selected follows a layered structure. Patrols reinforce ground presence in the main operational regions and lanes, while drone sorties expand coverage over nearby hotspot regions and sensors provide persistent observation at camps, waterholes, and a small number of strategic interior sites. Together, this forms the hotspot-focused strategy, which, as aforementioned, is this paper's key recommendation for resource allocation in Etosha. Full placement tables are given in Appendix F.

§10. Staffing and Protection Over Time

First, we must distinguish between two types of staffing:

- **active staff**, referring to the number of people who are working at a given time in direct operational support
- **payroll staff**, referring to the (larger) total headcount needed to keep that active presence running continuously (24 hours a day, 7 days a week).

For a package with B_{patrol} patrol routes, B_{drone} drone sorties, and B_{sensor} sensor sites, the number of active staff needed would equal

$$N_{\text{active}} = N_{\text{base}} + 6B_{\text{patrol}} + 2B_{\text{drone}} + \left\lceil \frac{B_{\text{sensor}}}{4} \right\rceil,$$

with $N_{\text{base}} = 12$ representing the operational personnel who are always needed.

We require a conversion factor between active staff and payroll staff, which we create using a simple coverage argument: if one employee contributes about 7 effective working hours per day, once rest, time off, and rotation are taken into account, then one continuously staffed role requires approximately $\frac{24}{7} \approx 3.5$ and hence

$$N_{\text{payroll}} \approx 3.5\, N_{\text{active}}.$$

> **Assumption** (7)**.** Note that 3.5 is simply a planning level simplification. It does not account for variations in work hours due to overtime work, holidays, weekends, or corporate differences. Of course, there are not exactly 3.5 shifts; this simply means that one active position sustained around the clock takes about three and a half people.

For the baseline combined package,

$$N_{\text{active}} = 12 + 6(10) + 2(4) + \left\lceil \frac{12}{4} \right\rceil = 83, \implies N_{\text{payroll}} \approx 3.5(83) \approx 291.$$

So the baseline package meets the 295 staff ceiling well. This attests to the validity of our calibration assumptions.

Beyond this baseline, in order to estimate staffing needs, we rerun the optimization at larger patrol capacities and record the smallest active field strength required to reach a target protection level. The technical derivation of this staffing curve is deferred to Appendix D1, and the full numerical outputs are deferred to Appendix F.

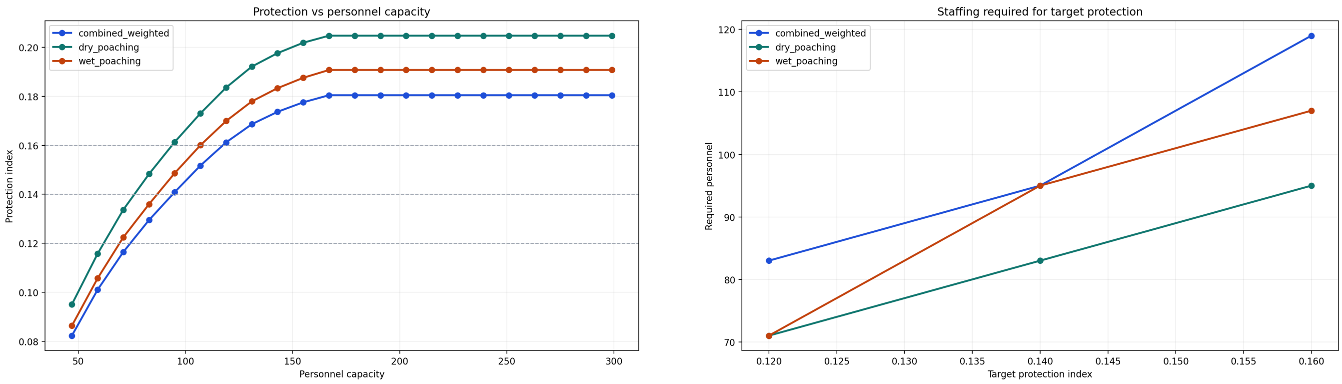

(a) Protection curve (b) Target requirements

Figure 9: Staffing need. Protection rises with active capacity, then gradually flattens.

Table 2: Minimum active staff counts for selected targets

Scenario	PI = 0.12	PI = 0.14	PI = 0.16
Dry poaching	71	83	95
Wet poaching	71	95	107
Combined weighted	83	95	119

In the combined scenario, reaching benchmarks of $PI = 0.12$, 0.14, and 0.16 requires 83, 95, and 119 active staff respectively, which correspond to approximately 291, 333, and 417 payroll staff (using the 3.5× conversion). The first benchmark is feasible under the current ceiling; the next two are not. The curve on the left also flattens near 167 active staff, or about 585 payroll staff, indicating that under the current candidate libraries, most of the best deployment opportunities have already been captured by that point.

Protection over time adds a second staffing lesson. The MILP chooses one strong strategic package for one calibrated scenario, but the landscape does not remain fixed after that choice is made. We therefore propagate the deployed system through a deterministic update layer in which unresolved loss raises future hazard, sustained monitoring suppresses it, and fire damage in the combined scenario temporarily reduces consequence before recovery. The update equations are deferred to Appendix D2.

> **Assumption** (8)**. Static allocation with dynamic refresh.** The allocation decision is solved as a scenario-static MILP, while temporal change is represented afterward in a deterministic update module. This keeps the strategic solve tractable while still allowing protection needs to evolve over time.

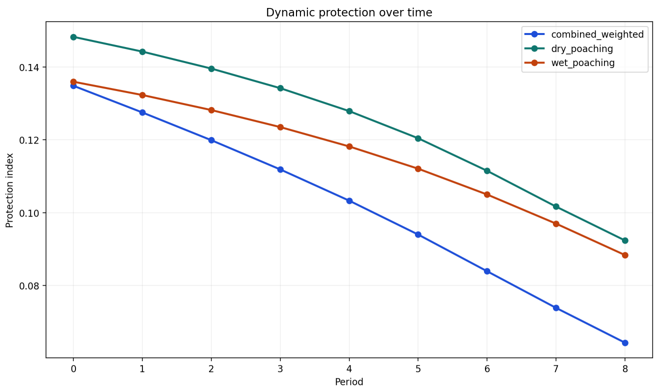

(a) Protection over time

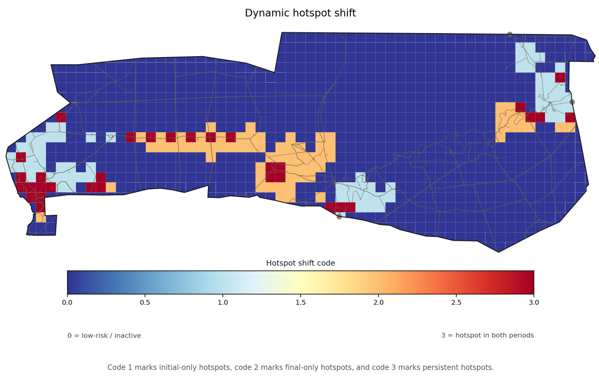

(b) Hotspot shift

Figure 10: Protection over time. If the baseline package is left unchanged, persistent pressure gradually shifts hotspots and weakens protection.

In the combined scenario, protection begins near 0.135 at period 0 and falls to about 0.064 by period 8 if the original deployment is left unchanged. This is best interpreted as a management result rather than a literal forecast: a good baseline plan still needs to be rerun as pressure accumulates and hotspots shift.

§11. Stress Tests, Uncertainty, and Equity

We believe that a good recommendation should survive reasonable changes in assumptions and resources. We therefore test the baseline in three ways: across multiple scenarios, under direct stress cases, and under centered Monte Carlo perturbations. The exact implementation is deferred to Appendix F.

The three main operating scenarios are dry-season poaching, wet-season poaching, and the combined weighted case. The combined case is stressed by a 30% staff shortfall, loss of two drone sorties, dropout of four sensors, a 30% slowdown in response, a poaching surge, and a more severe fire year.

Table 3: Combined-scenario stress ranking

Case	PI	Relative change from baseline
Baseline	0.129	0.0%
Poaching surge	0.131	+1.4%
Sensor dropout	0.127	-2.0%
Drone loss	0.125	-3.8%
Fire year	0.122	-5.4%
Slow response	0.118	-8.8%
30% staff shortfall	0.109	-15.8%

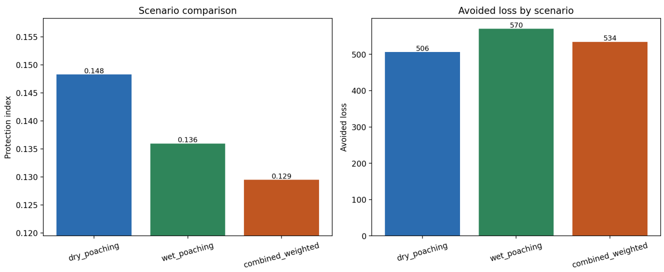

(a) Deterministic scenarios

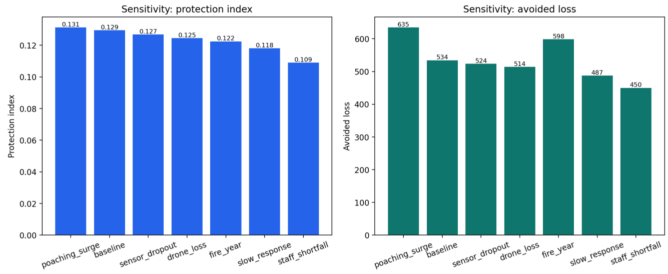

(b) Combined-scenario stress tests

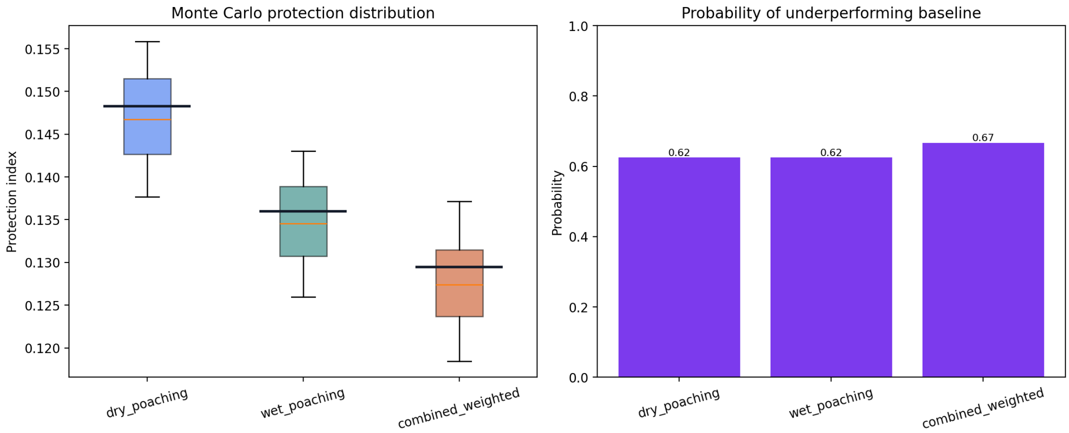

(c) Monte Carlo robustness

Figure 11: Stress and uncertainty. The baseline recommendation is directionally stable, but it is much more sensitive to loss in response capacity than to moderate losses of drones or sensors.

The message is clear in that moderate losses of sensors or drones hurt performance, but slower response and staff shortfall hurt more. That ranking follows directly from the objective logically as monitoring has value only when it can still be converted into timely intervention. In particular we would like to note that the slight rise in PI under "poaching surge" does not mean worsening poaching helps; it means the added pressure shifts toward cells that are already better covered, so the fraction of loss avoided increases slightly even while total pressure worsens.

The Monte Carlo results say the same in probabilistic form. In the combined scenario, the deterministic baseline is $PI = 0.129$, while the Monte Carlo mean is 0.128 with standard deviation about 0.0051 and a 5th–95th percentile interval of roughly 0.120–0.134. The solution is therefore reasonably stable, but uncertainty is asymmetric: degraded response hurts more sharply than equally sized improvements help.

Validation produces one reassuring conclusion and one policy trade-off. The reassuring conclusion is that the strategic pattern is stable: across grid sizes of 4, 5, and 7.500 km, and under moderate perturbations to the waterhole and response layers, the combined protection index stays in a very narrow band near 0.129–0.134. We note the slightly higher values of PI in other sizes should not be interpreted as a superior deployment, but simply modest discretization sensitivity. The harder conclusion is that the unconstrained hotspot-first optimum is not spatially even.

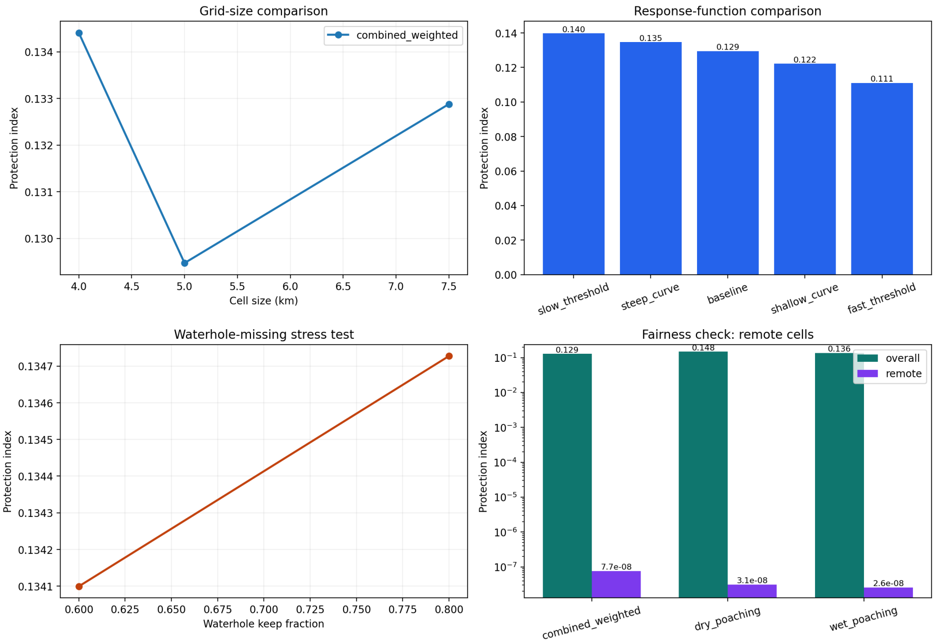

Figure 12: Validation summary. Moderate structural perturbations preserve the main strategic ranking, while the remote tail remains under-protected in the unconstrained optimum.

In the combined scenario, the most remote 10% of cells have:

- remote protection index essentially equal to zero,
- mean monitoring only 0.028 versus 0.828 park-wide,
- mean response time 34.63 hours versus 5.23 hours park-wide.

This is the main trade-off in the Etosha baseline. The strategy is efficient because it concentrates effort where monitoring can still become action. It is inequitable because the hardest-to-reach tail of the park receives very little effective service.

What the Etosha results say

Etosha should be protected by a hotspot-first baseline strategy. Under the current staffing ceiling, the best feasible use of scarce patrols, drones, and sensors is to concentrate them where consequence, threat, and reachability overlap. If leadership wants a higher service floor in remote sectors, that can be imposed, but it will come with a measurable loss of total efficiency.

§ Discussion and Generalization

§1. Strengths

1. In our model, the concept of protection is clearly and consistently defined as avoided ecological loss. This strengthens our paper by guiding our analysis of threat, value, and resource allocation.
2. In our model, all maps and layers are realistic to Etosha, since road, gate, camp, waterhole, and terrain data were all directly taken from open-source datasets, as discussed in Appendix A. Furthermore, this is highly generalizable since such data for other parks can be similarly

accessed.

3. Our model connects theoretical ecological importance to operational realities. Value, threat, detection, and response are synthesized under one framework so that conservation can take all these factors into account.
4. The results produced by the model are easy to interpret for decision-maker use. Recommended deployment locations, a staffing curve, and clear scenario comparisons are all provided so the outputs can support practical planning.

§2. Limitations

1. While the model is strategic, it is not fully operational. While we do show how resources should be optimally allocated, we do not schedule personnel in real-time or make shift-by-shift dispatch decisions.
2. Several operational details are simplified. For instance, staff are treated as homogeneous, equipment reliability is assumed rather than further modeled, and intervention success is taken for granted.
3. Adversaries are assumed to be boundedly rational, which could ignore game-theoretic influence on decision-making. For one, poachers are not modeled to be strategic agents that take the park's protection mechanisms into full account.

Each of these limitations point to solid future work–consult Appendix H.

§3. Generalization Framework

Our modeling does not end with Etosha; we also present a framework that can be applied to other national parks and ecological reserves. The critical idea that we discovered is that the optimization principles and mechanisms (maximizing PI, using MILP, etc.) can feasibly be transferred, but new calibration metrics and local geographical data will need to be introduced.

To generalize to another park, the same workflow will use a portable geospatial pipeline to automatically pull in global data and calibrate according to local ecology, reconstructing the park's basic geographical structure. Once a new park's map is given, we can rebuild much of the spatial backbone from public sources: OpenStreetMap for roads and access features, OpenTopography for deriving slope and visibility, and NASA FIRMS together with NASA POWER to provide fire detection and fire-weather data. In that sense, the grid layout, road graph, terrain layer, and fire hazard maps are flexible with respect to location, rather than strictly fixed to Etosha.

However, some aspects of the model are non-transferrable, and those are just as significant. Etosha's ecological value layer is heavily linked to a unique waterhole environment, which is not universal. Hence, a complete generalization, in addition to drawing new boundaries, must also adjust to new key ecological features, threat sources, operating bases, and more.

Table 4: Generalization steps

Component	Necessary changes
Loss-reduction objective and MILP implementation	• unchanged; the principle of loss reduction, the decision variables, and the detection surrogate remain the same
Hybrid grid–graph layout	• regenerated automatically from the new boundary, road network, and DEM
Ecological value layer	• replaced by park-specific species, habitat, refuge, or visitor-conflict attractors
Threat layer	• replaced by park-specific risks (e.g. poaching, tourism pressure, flooding, extreme weather)
Response-time logic	• unchanged in form, but recomputed from local slopes, and area access
Fire-context workflow	• same ingestion pipeline, but recalibrated to account for local weather and landscape
Bases and routes	• rebuilt from local camps, gates, trailheads, etc.
Budgets, staffing, and fairness targets	• adapted to local policy and resource constraints

§4. Two Adaptations to Other Protected Areas

Etosha is the main park that we solved completely; however, we deliberately generalize to 2 contrasting national parks to show how the same architecture behaves when ecological and operational circumstances change drastically. Namely, we consider Yellowstone National Park and Kaziranga National Park.

Table 5: Two evidence-based adaptation summary

Park	Notable differences	Expected model changes
Yellowstone	High tourism, extensive wetlands and river systems, ecologically important wildfire, and seasonal road closures (National Park Service, 2025a; National Park Service, 2025b; National Park Service, 2025c; *Fire*, 2026; *Park Roads*, 2026)	Replace scarce-water attractors by habitat and visitor-conflict layers; increase wildfire weight; impose snow-season road closures; shift monitoring toward entrances, campgrounds, major junctions, and wildlife-viewing corridors
Kaziranga	Dynamic Brahmaputra floodplain, annual flooding, globally important rhino habitat, and officially recognized animal corridors toward higher ground (UNESCO World Heritage Centre, n.d.; Government of Assam, Finance Department, 2022; District Disaster Management Authority (DDMA), Golaghat, 2025; Government of Assam, 2021)	Replace Etosha's waterhole value field by rhino habitat, flood refuge, and corridor layers; make the mobility graph flood-sensitive; prioritize chokepoints, crossings, highlands, and anti-poaching bases; allow seasonal redeployment when camps or routes are inundated

Unlike how risk in Etosha is clearly concentrated around waterholes, in Yellowstone, risk is instead influenced by high visitor pressure and intense seasonal road constraints. On the other hand, in Kaziranga, poaching remains a highly relevant issue, yet the operational picture changes sharply because movement is shaped by flood-sensitive passageways, refuge areas, and seasonal inundation rather than by a relatively stable road network. These two parks were chosen precisely because they put the effectiveness of the model under scrutiny from different angles. On one hand, Yellowstone tests whether the framework will still work when ecological value is no longer centered around areas like waterholes; meanwhile, on the other hand, Kaziranga can check whether the modeling approach still holds when park access and mobility fluctuate. Appendix I gives the empirical rerun.

§5. Adaptable Deployment Methods

The correct deployment path will begin with calibration. Clearly, conservation in a new park should not utilize Etosha's ecological weights, access assumptions, or asset priorities and then be mechanically optimized. In particular, metrics like local ecological value determinants, threat sources, planning budgets, human access structure, and seasonal restrictions should be changed. Some required inputs would be the new ecological habitat features, human access layers, park budget and staffing constraints, and closure or weather data.

The deployment itself should shift with the park. In Yellowstone, protection would likely lean toward entrances, visitor attractions, and roads that are seasonally open. In Kaziranga, it would likely shift to crossings, refuge areas, and flood-proof access points. As an example, the empirical reruns in Appendix I show this process carried outin practice.

§ Conclusion

The challenge presented to us was to create a generalized protection strategy for national wildlife reserves. Starting with Etosha National Park, we developed an adaptable optimization model that measured protection as avoided ecological loss rather than map coverage. By integrating ecological value, threat intensity, detection capacity, and response times, the model generated a hotspot-first deployment strategy consisting of 10 patrol routes, 4 drone sorties, and 12 sensor sites. This baseline package requires 83 active personnel, or 291 total staff under full operations, placing it just below the park's 295 person limit. Built on terrain-aware travel times, seasonal threat mapping, and a variety of other ecological and operational factors, our mixed integer linear program model offers both a realistic plan for Etosha and a transferable structure for other parks. Its preliminary transfer to Yellowstone and Kaziranga suggests that while each reserve requires local recalibration, the overall modeling approach can be applied far beyond a single case.

§ References

Copernicus Programme. (n.d.). *Copernicus Digital Elevation Model (COP-DEM) [Data set]*. Copernicus Data Space Ecosystem. Retrieved March 19, 2026, from https://dataspace.copernicus.eu/explore-data/data-collections/copernicus-contributing-missions/collections-description/COP-DEM

District Disaster Management Authority (DDMA), Golaghat. (2025). *District Disaster Management Plan, Golaghat 2025–26*. Assam State Disaster Management Authority (ASDMA). https://asdma.assam.gov.in/sites/default/files/swf_utility_folder/departments/asdma_revenue_uneecopscloud_com_oid_70/menu/document/golaghat_2025_26.pdf

Fire. (2026, March 12). U.S. National Park Service. Retrieved March 19, 2026, from https://www.nps.gov/yell/learn/nature/fire.htm

Ghoddousi, A., Van Cayzeele, C., Negahdar, P., Soofi, M., Kh. Hamidi, A., Bleyhl, B., Fandos, G., Khorozyan, I., Waltert, M., & Kuemmerle, T. (2022). Understanding spatial patterns of poaching pressure using ranger logbook data to optimize future patrolling strategies. *Ecological Applications, 32*(5). https://doi.org/10.1002/eap.2601

Government of Assam. (2021). *Flood Memorandum to the Government of India: Assam Floods 2021*. Government of Assam. https://asdma.assam.gov.in/sites/default/files/swf_utility_folder/departments/asdma_revenue_uneecopscloud_com_oid_70/menu/document/assam_flood_memorandum_2021.pdf

Government of Assam. (2022, April 1). *Notification on Delineation of Nine Animal Corridors Connecting Karbi Anglong Hills and Kaziranga National Park*. Government of Assam. https://dpns.assam.gov.in/sites/default/files/swf_utility_folder/departments/directorate_printing_uneecopscloud_com_oid_3/menu/document/no._253_frw_7-2022-01_dated_01-04-22_part-i.pdf

Government of Assam, Finance Department. (2022, April 1). *No. 253-FRW/7/2022/1 (Part I)*. Government of Assam. https://dpns.assam.gov.in/sites/default/files/swf_utility_folder/departments/directorate_printing_uneecopscloud_com_oid_3/menu/document/no._253_frw_7-2022-01_dated_01-04-22_part-i.pdf

Government of Namibia. (n.d.). *Etosha National Park*. Ministry of Environment, Forestry and Tourism. Retrieved March 19, 2026, from https://www.meft.gov.na/national-parks/etosha-national-park/217

Haklay, M., & Weber, P. (2008). OpenStreetMap: User-generated street maps. *IEEE Pervasive Computing, 7*(4), 12–18. https://doi.org/10.1109/mprv.2008.80

Map of Etosha National Park. (n.d.). Etosha National Park. Retrieved March 19, 2026, from https://etoshanationalpark.co.za/etosha-map/

NASA. (n.d.-a). *Fire Information for Resource Management System (FIRMS) active fire data [Data set]*. NASA Earthdata. Retrieved March 19, 2026, from https://firms.modaps.eosdis.nasa.gov/active_fire

NASA. (n.d.-b). *Prediction of Worldwide Energy Resources (POWER) data [Data set]*. NASA POWER Project. Retrieved March 19, 2026, from https://power.larc.nasa.gov/

NASA. (2025, March 31). *Prediction of Worldwide Energy Resources (POWER) [Data set]*. NASA Open Data Portal. Retrieved March 19, 2026, from https://data.nasa.gov/dataset/prediction-of-worldwide-energy-resources-power

National Park Service. (2025a, April 18). *Park facts*. NPS. Retrieved March 19, 2026, from https://www.nps.gov/yell/planyourvisit/parkfacts.htm

National Park Service. (2025b, December 2). *Visitor use management*. NPS. Retrieved March 19, 2026, from https://www.nps.gov/yell/learn/management/visitor-use.htm

National Park Service. (2025c, April 18). *Wetlands*. NPS. Retrieved March 19, 2026, from https://www.nps.gov/yell/learn/nature/wetlands.htm

Ndachena, N., Toorians, M. E., Farrell, M. J., Govender, D., & Davies, T. J. (2025). Effect of drought on wildlife activity at artificial waterholes. *Biological Conservation, 310*, 111370. https://doi.org/10.1016/j.biocon.2025.111370

OpenStreetMap contributors. (n.d.). *About OpenStreetMap*. OpenStreetMap Foundation. Retrieved March 19, 2026, from https://www.openstreetmap.org/about/eng

OpenTopography. (2021). *Global DEM (OTSDEM.032021.4326.3) [Data set]*. OpenTopography. Retrieved March 19, 2026, from https://portal.opentopography.org/raster?opentopoID=OTSDEM.032021.4326.3

Park roads. (2026, March 19). NPS. Retrieved March 19, 2026, from https://www.nps.gov/yell/planyourvisit/parkroads.htm

Protected Areas (WDPA). (n.d.). Protected Planet. Retrieved March 19, 2026, from https://www.protectedplanet.net/en/thematic-areas/wdpa?tab=WDPA

Pulido-Chadid, K., Virtanen, E., & Geldmann, J. (2023). How effective are protected areas for reducing threats to biodiversity? A systematic review protocol. *Environmental Evidence, 12*(1). https://doi.org/10.1186/s13750-023-00311-4

Pyomo developers. (n.d.). *About Pyomo*. Pyomo. Retrieved March 19, 2026, from https://www.pyomo.org/about

Tobler, W. R. (1993, February). *Three presentations on geographical analysis and modeling*. Lukatela. Retrieved March 19, 2026, from https://www.lukatela.com/hrvoje/papers/tobler93.html

Turner, W. C., Périquet, S., Goelst, C. E., Vera, K. B., Cameron, E. Z., Alexander, K. A., Belant, J. L., Cloete, C. C., du Preez, P., Getz, W. M., Hetem, R. S., Kamath, P. L., Kasaona, M. K., Mackenzie, M., Mendelsohn, J., Mfune, J. K., Muntifering, J. R., Portas, R., Scott, H. A., ... Kilian, J. W. (2022). Africa's drylands in a changing world: Challenges for wildlife conservation under climate and land-use changes in the greater Etosha landscape. *Global Ecology and Conservation, 38*, e02221. https://doi.org/10.1016/j.gecco.2022.e02221

UNESCO World Heritage Centre. (n.d.). *Kaziranga National Park*. UNESCO World Heritage Convention. Retrieved March 19, 2026, from https://whc.unesco.org/en/list/337/

van Doormaal, N., Lemieux, A. M., Ruiter, S., Allin, P. M. R. R., & Spencer, C. R. (2021). Detecting wildlife poaching: A rigorous method for comparing patrol strategies using an experimental design. *Oryx, 56*(4), 572–580. https://doi.org/10.1017/S0030605320001301

§ AI Use Report

In addition to the four main uses below, in which the outputs and prompts are displayed, we used AI sporadically throughout the paper to give limited wording and formatting assistance, smoothening some sentences and summarizing our thoughts. In addition, AI gave us some insights on how to structure and format the appendices.

For most AI use, we used Perplexity AI (GPT 5.4) for its live, citation-backed web search.

Use 1: Map layouts and coding structures' support

Model: Perplexity AI GPT 5.4

Prompt

```
I'm writing an modelling paper about protecting wildlife in Etosha National Park, and I'm building all of the maps for the paper in Python. I need a clean, consistent layout for the full map set. The figures I want included: an Etosha boundary map, camps map, gates map, roads map, waterholes map, a plain grid map, dry and wet consequence maps, dry and wet hazard maps, a response-time map, a selected deployment map showing patrol routes, drone anchors, and sensor sites, and a few context establishment maps. Can you help me design a good coding structure for these maps so they all look like they belong in the same paper? Please, on top, suggest/tip how to organize the plotting code, keep fonts/colors/line widths consistent, choose figure sizes, export files cleanly, and decide which maps should be standalone versus multi-panel.
```

Output

```
Yes. The cleanest way to make the map set feel like one paper is to treat it like a small plotting system, not a collection of one-off scripts.

I would organize it around three layers:
1. one shared style file for fonts, colors, line widths, marker sizes, legends, scale bars, and export settings;
2. one reusable map-building pipeline with functions like draw_base_map, plot_roads, plot_waterholes, plot_grid, and plot_patrol_routes;
3. a small set of fixed figure templates so every map does not invent its own layout.

A practical structure would be:
maps/
  src/
    styles.py
    layers.py
    layout.py
    map_specs.py
    export.py
    build_all.py
  outputs/
    figures/

For the style system, I would put things like boundary color, road color, waterhole color, patrol/drone/sensor colors, font sizes, and linewidth tokens in one MapStyle object. That way you can change the whole paper's visual language in one place instead of hunting through every script.
```

```

For figure structure, I would use only a few templates:
- full-width standalone for response-time and selected deployment;
- paired 1x2 panels for dry/wet consequence and dry/wet hazard;
- compact multi-panel plates for reference layers such as camps, gates, roads, waterholes, and grid.

The biggest consistency rules are:
- use the same extent for all full-park analytical maps;
- keep dry and wet versions on the same color scale;
- keep legends, north arrows, and scale bars stylistically identical;
- use one export function and one naming scheme for all outputs.

If you want a concrete starting point, build a MapSpec object per figure, then let build_all.py loop through those specs and call shared layer functions in order. That will make the map suite much easier to revise as the paper evolves.
```

Use 2: External API integration

Model: Perplexity AI GPT 5.4

Prompt

```
i plan to hook external APIs such as OpenTopography and NASA FIRMS. how can I do that?
```

Output

```
Yes. The clean way to do this is to treat OpenTopography and FIRMS as data providers, not as part of the plotting code itself.

I would split the workflow into four stages:
1. fetch raw data from the API;
2. normalize it into your own GIS formats;
3. cache the processed result locally;
4. make the map scripts read only from local files.

That means your figure code never depends on a live request.

A good project split would be:
src/
  api/
    opentopography.py
    firms.py
  processing/
    dem.py
    fire.py
  build_data.py
  build_maps.py

OpenTopography should mainly feed DEM-based products such as terrain, slope, hillshade, and possibly visibility or movement-cost layers. FIRMS should be handled as a fire-event source that gets converted into stable local layers such as filtered point files, seasonal fire summaries, or fire-density rasters.

```

```
The most important implementation rules are:
- store API keys in environment variables, not in the repo;
- define the Etosha AOI once and use one request CRS, then reproject after download;
- cache raw files so repeated runs do not hammer the APIs;
- save metadata like source, date range, product, and AOI in sidecar files;
- have one command for rebuilding data and a separate command for rebuilding maps.

In short: use the APIs to build your processed layers, but never let your final paper figures depend on live API calls. That is the most reproducible and contest-safe pattern.
```

Use 3: Park feature and transfering research

Model: Perplexity AI GPT 5.4

Prompt

```
I'm using Etosha as the main case study in my paper, and I also want to discuss Yellowstone and Kaziranga as transfer cases. Can you help me summarize the park features that matter for a protection model in each park? I'm especially interested in geography, hydrology, roads and access, seasonal effects, wildlife concentration patterns, and operational constraints. Please focus on the kinds of features that would affect patrols, drones, sensors, response time, and spatial risk.
```

Output

```
Yes. The three parks are useful because they stress very different parts of a protection model.

Etosha is the clearest waterpoint-and-road-network case. Wildlife concentration is strongly shaped by waterholes, especially in the dry season, and access is comparatively road-based. That makes Etosha well suited to a node-and-corridor model in which patrols, sensors, and drones cluster around waterholes, camps, gates, road corridors, and pan-edge movement zones.

Yellowstone is not a scarce-water case. It is better understood as a seasonal road-access and valley-concentration case. Hydrology is widespread, topography is much more complex, and winter road closures matter a great deal. A transferred model would need to replace Etosha's waterhole logic with habitat, valley, migration, and seasonal accessibility structure, especially if response times change sharply between summer and winter.

Kaziranga is the strongest floodplain-corridor case. Hydrology is not just background geography there; it drives the problem. Flooding changes wildlife concentration, corridor use, road-crossing conflict, patrol feasibility, sensor survivability, and response-time geometry. A transferred model would likely need separate normal and flood-season scenarios, or an explicit flood-state variable.

So the simplest cross-case story is:
- Etosha: scarcity and concentration around waterholes;
- Yellowstone: seasonal access, valleys, and topography;
- Kaziranga: flooding, refuge corridors, and crossing hotspots.
```

```

That is a strong transferability narrative because it shows that the optimization architecture can stay, while the ecological and operational calibration has to change substantially from park to park.
```

Use 4: LaTeX-ing figure formats

Model: Gemini 3.1 Pro

Prompt

```
I'm writing my modelliong paper in overleaf latex and I want one figure made of five subfigures: three across the top row and two across the bottom row, with one shared caption for the whole figure and separate small captions for each panel. Can you show me a clean LaTeX template for that layout using subfigures, and make it suitable for a formal contest paper?
```

Output

```
Yes. A clean contest-style way to do this is to use subcaption with one figure environment, three panels across the top row, two on the bottom row, and one shared caption.

I would use widths around 0.31\textwidth for the top row and 0.47\textwidth for the bottom row. That usually balances well without making the rows look crowded.

The basic structure is:
\begin{figure}[t]
    \centering
    \begin{subfigure}[t]{0.31\textwidth}
        \centering
        \includegraphics[width=\linewidth]{panel_a.pdf}
        \caption{Boundary}
    \end{subfigure}
    \hfill
    \begin{subfigure}[t]{0.31\textwidth}
        \centering
        \includegraphics[width=\linewidth]{panel_b.pdf}
        \caption{Camps}
    \end{subfigure}
    \hfill
    \begin{subfigure}[t]{0.31\textwidth}
        \centering
        \includegraphics[width=\linewidth]{panel_c.pdf}
        \caption{Gates}
    \end{subfigure}

    \vspace{0.8em}

    \begin{subfigure}[t]{0.47\textwidth}
        \centering
        \includegraphics[width=\linewidth]{panel_d.pdf}
        \caption{Roads}
    \end{subfigure}
```

```
    \hfill
    \begin{subfigure}[t]{0.47\textwidth}
        \centering
        \includegraphics[width=\linewidth]{panel_e.pdf}
        \caption{Waterholes}
    \end{subfigure}

    \caption{Reference maps for Etosha National Park used in the protection model.}
\end{figure}

The main formatting advice is:
- keep the subcaptions short;
- put the real explanation in the main caption;
- use visually matched panels;
- prefer vector files for maps when possible.

That layout is especially good for your boundary/camps/gates/roads/waterholes reference figure because it saves space and still reads cleanly.
```

In all four main-use cases, the prompts were kept saved and the exact AI responses were recorded. Final modeling choices, code, figures, and interpretations were all thoroughly checked by the team before any inclusion in the paper.

§ Appendix

§1. Appendix A: Data Inventory and External Resources

This appendix collects the data resources, public APIs, and acquisition routes used in the present Etosha analysis.

Geometric Foundations

Four physical entities dominate the present case study:

- **Waterholes**, which concentrate dry-season wildlife value and consequently influence poaching hazard;
- **Roads,** which shape response-times (especially to wildfires);
- **Camps and gates,** which act as response bases and access choke-points;
- **Terrain,** which modifies travel time and line of sight.

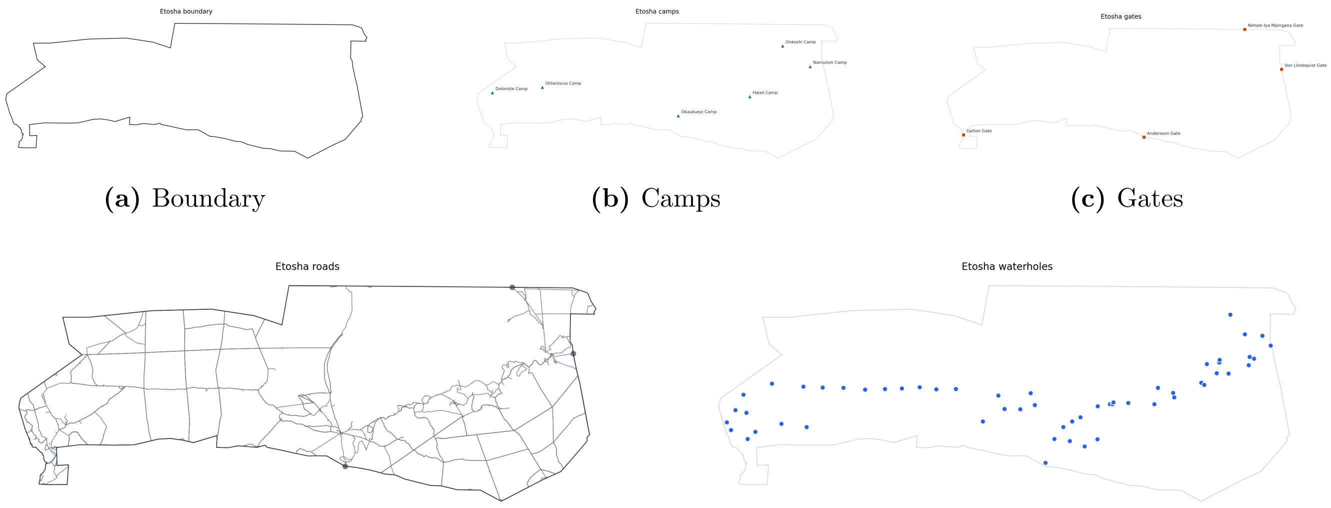

(a) Boundary **(b)** Camps **(c)** Gates

(d) Roads **(e)** Waterholes

Figure 13: Etosha context layers, the geometric foundations from which the value field, threat field, movement graph, and asset libraries are built.

Table 6: Data inventory and external resources for the Etosha analysis

Layer or resource	Source	Present role	Notes
Park boundary	OpenStreetMap (OpenStreetMap contributors, n.d.)	clipping, boundary distance, local coordinate origin	stored as processed boundary geometry
Road network	OpenStreetMap (OpenStreetMap contributors, n.d.)	road graph, access covariates, patrol library	acquired from OSM and cached
Camps, gates, waterholes	Etosha reference overlay aligned to official map (*Map of Etosha*, n.d.)	response bases, ecological attractors, sensor sites	56 waterholes, 6 camps, 4 gates in the current run

Layer or resource	Source	Present role	Notes
Terrain raster	OpenTopography global DEM service (OpenTopography, 2021)	elevation, slope, off-road speed, line of sight	current Etosha run uses `dem.tif` from the terrain fetch pipeline
Active and recent fire detections	NASA FIRMS (NASA, n.d.-a)	recent and historical fire density features	recent detections and history are cached locally
Fire-weather context	NASA POWER (NASA, 2025)	dryness and wind indices	365 daily rows in the current run
Official park map	Etosha National Park map (*Map of Etosha*, n.d.)	spatial verification of camps, gates, and waterholes	used in node auditing and narrative checks
IMMC problem brief	IMMC prompt	policy framing, staff ceiling, problem requirements	defines 295-person staff ceiling

Table 7: Current processed Etosha inventory

Quantity	Value	Comment
Clipped grid cells	996	5.000 km cells
Waterholes	56	authoritative Etosha reference set
Camps	6	active bases
Gates	4	access nodes and response bases
Patrol routes in library	24	camp-cluster, gate-screen, perimeter-screen
Drone sorties in library	12	candidate sweeps
Sensor candidates	26	camps, gates, waterholes, strategic sites
Monte Carlo trials	24	centered robustness runs
Dynamic periods	9	period 0 through period 8

Note

We used open source tools for geospatial acquisition, optimization, and reproducibility. For completeness, the present implementation uses an automated OpenStreetMap acquisition layer, an open-source MILP modeling layer (Pyomo developers, n.d.), and a standard MILP solver.

§2. Appendix B: Etosha Design Notes

In simple terms we try to explain the design logic of our Etosha model, in this appendix. The underlying principle is that protection matters most where ecological value, external threats, and fast response all meet in the same place. The model works like a map of pressure points, and we optimize in order to emphasize the parts of the park where intervention can save the most value.

Waterholes dictate the value layer

In Etosha, waterholes are wildlife attractors – as after all, all animals need water to survive. Trivially, the more wildlife present in a region, the higher the ecological stakes. For that reason, the value layer gives strong weight to waterhole proximity.

Separating value from threat

The model treats ecological consequence and human threat as two separate layers simply for clarity. One layer shows where loss would hurt most, while another shows where threats are most likely to appear. Not only do their combination gives a better visualization of the situation, having a high value of one does not imply the same in the other.

Why roads are so important

Roads shape the park in two ways. One, they help ranger teams move quickly. Second, they make approach and escape easier for intruders. In that sense, roads encompass movement in both directions/sides, which the model includes to reflect the real strategy.

Response time defines contestability

Detection helps most when help can arrive quickly. That is why the model focuses on travel time instead of map distance alone. Cells with fast access remain much more contestable compared to slower cells.

Why use patrols, drones, and sensors?

We utilize all three assets because each fulfills a unique role. Patrols provide visible presence and direct intervention capability. Drones extend reach across large areas at a time. Sensors offer persistent monitoring at important points. Together, they create a complete protection system for the park.

Grid and candidate libraries

The model constructs a library of feasible patrol routes, drone sorties, and sensor sites, which keeps the final deployment strategy easy to interpret.

How to read the loss scores

The loss values are best understood when compared across different cells, since they operate on a relative scale. They allow us prioritize locations, compare deployments, as well as show which allocation strategy produces the greatest reduction in expected loss. Again, this is central to our model.

Design takeaway

The Etosha model concentrates protection on the overlapping of the three forces of (1) high ecological value, (2) credible human pressure, and (3) response that arrives quickly enough to matter.

§3. Appendix C: Full Notation Table

Table 8: Full notation table

Symbol	Definition	Type
G	Set of analysis cells in the clipped park grid	Index set
$\mathcal{R}$	Candidate patrol routes	Index set
$\mathcal{D}$	Candidate drone sorties	Index set
$\mathcal{S}$	Candidate sensor sites	Index set
T	Dynamic periods in the update module	Ordered index set
H	Hotspot-cell set selected from the upper baseline-loss tail	Index set
Q	Remote-cell set used in the optional fairness analysis	Index set
$g, r, d, s, t, k, \omega, \sigma, \ell$	Cell, route, sortie, site, time, threat-type, scenario, season, and piecewise-envelope segment indices	Indices
A_g	Area of cell g after clipping	Parameter, positive real
c_g	Centroid of cell g	Parameter, geometric point
a_d	Anchor point of drone sortie d	Parameter, geometric point
Γ_r	Polyline geometry of patrol route r	Parameter, geometric object
$d_{g,W}$	Distance from cell g to nearest waterhole	Parameter, nonnegative real (km)
$d_{g,\text{gate}}$	Distance from cell g to nearest gate	Parameter, nonnegative real (km)
$d_{g,\text{road}}$	Distance from cell g to nearest road	Parameter, nonnegative real (km)
$d_{g,\text{camp}}$	Distance from cell g to nearest camp	Parameter, nonnegative real (km)
$d_{g,\text{fence}}$	Distance from cell g to park boundary / fence	Parameter, nonnegative real (km)
$d_{g,\text{ctr}}$	Distance from cell g to the park-center surrogate used in the fire field	Parameter, nonnegative real (km)
$q^{(r)}_{g,\sigma}$	Rhino-use field in season σ	Parameter, $[0, 1]$
$q^{(e)}_{g,\sigma}$	Elephant-use field in season σ	Parameter, $[0, 1]$
$q^{(h)}_{g,\sigma}$	Herbivore-use field in season σ	Parameter, $[0, 1]$
$\omega_r, \omega_e, \omega_h$	Species weights used in the consequence layer for rhino, elephant, and herbivore components	Parameters, nonnegative reals
B_g	Structural biomass surrogate used in fire consequence and hazard layers	Parameter, $[0, 1]$
H_g	Normalized historical fire density from FIRMS detections	Parameter, $[0, 1]$
J_g	Normalized recent fire density from FIRMS detections	Parameter, $[0, 1]$
X_g	Normalized human-access pressure field from roads, gates, and fence proximity	Parameter, $[0, 1]$

Symbol	Definition	Type
S_g^{slope}	Normalized slope-related wildfire spread factor	Parameter, $[0, 1]$
D_t	Weather dryness index at time t	Parameter, $[0, 1]$
W_t	Weather wind index at time t	Parameter, $[0, 1.5]$
C_g	Ecological consequence score at cell g in generic form	Parameter/state, $[0, 1]$
$C_{g,\sigma}$	Ecological consequence score at cell g in seasonal form	Parameter/state, $[0, 1]$
$C_{g,t}$	Ecological consequence score at cell g in dynamic time-indexed form	Parameter/state, $[0, 1]$
λ_g	Threat intensity score at cell g in generic form	Parameter/state, $[0, 1]$
$\lambda_{g,\omega}$	Threat intensity score at cell g under scenario ω	Parameter/state, $[0, 1]$
$\lambda_{g,t}$	Threat intensity score at cell g in dynamic time-indexed form	Parameter/state, $[0, 1]$
$\beta_r, \beta_G, \beta_R, \beta_F$	Dry/wet poaching-hazard weights on rhino, gate, road, and fence proximity terms	Parameters, non-negative reals
η_G, η_F, η_R	Intrusion-hazard weights on gate, fence, and road proximity terms	Parameters, non-negative reals
$\zeta_1, \ldots, \zeta_9$	Fire-hazard weights on biomass, center distance, fence, fire history, recent fire, access, dryness, wind, and slope terms	Parameters, non-negative reals
θ	Local terrain slope	Parameter, real
θ_{ij}	Terrain slope on road segment (i, j)	Parameter, real
$v_{\text{class}(ij)}$	Base travel speed assigned to the class of road segment (i, j)	Parameter, positive real
v_{ij}^{road}	Terrain-adjusted travel speed on road segment (i, j)	Derived parameter, positive real
$v^{\text{off}}(\theta)$	Off-road terminal travel speed as a function of slope θ	Function, positive real-valued
c	Road-slope penalty coefficient in the road-speed adjustment	Parameter, nonnegative real
$f_{\min}$	Minimum allowed road-speed factor under slope adjustment	Parameter, $(0, 1]$
τ_{bg}	Response time from base b to cell g	Parameter, positive real (h)
τ_g	Best response time over all bases for cell g	Parameter, positive real (h)
τ_0	Midpoint of the logistic response-success curve	Parameter, positive real (h)
κ	Steepness of the logistic response-success curve	Parameter, positive real
R_g	Conditional response success at cell g after detection	Parameter, $[0, 1]$
D_g	Detection success at cell g	Derived quantity, $[0, 1]$

Symbol	Definition	Type
P_g	Overall prevention success at cell g	Derived quantity, $[0, 1]$
ϕ_{rg}	Patrol influence of route r on cell g	Parameter, nonnegative real
ψ_{dg}	Drone influence of sortie d on cell g	Parameter, nonnegative real
χ_{sg}	Sensor influence of site s on cell g	Parameter, nonnegative real
$\alpha_P, \rho_P, \varepsilon_P$	Patrol influence scale, patrol distance-decay length, and patrol sparsity threshold	Parameters, nonnegative reals
$\alpha_D, R_d, \varepsilon_D$	Drone influence scale, sortie radius, and drone sparsity threshold	Parameters, nonnegative reals
$\chi_0, \chi_{\min}, \rho_S, R_S$	Sensor base detection, minimum detection floor, sensor detection-decay length, and sensor radius	Parameters, nonnegative reals
vis_{sg}	Visibility indicator equal to 1 if sensor site s has line of sight to cell g	Parameter, $\{0, 1\}$
M_g	Aggregate monitoring score induced by selected assets	Decision variable, nonnegative real
S_g	Bounded detection surrogate used in the MILP	Decision variable, $[0, 0.45]$
b_ℓ	Breakpoint of the piecewise detection envelope at segment index ℓ	Parameter, nonnegative real
v_ℓ	Anchor value of the piecewise detection envelope at segment index ℓ	Parameter, nonnegative real
m_ℓ	Slope of segment ℓ in the piecewise detection envelope	Parameter, nonnegative real
$u_\ell(\cdot)$	Affine upper-bounding segment in the piecewise detection envelope	Affine function
y_r	Equals 1 if patrol route r is selected	Decision variable, $\{0, 1\}$
z_d	Equals 1 if drone sortie d is selected	Decision variable, $\{0, 1\}$
x_s	Equals 1 if sensor site s is selected	Decision variable, $\{0, 1\}$
n_r	Staffing requirement attached to patrol route r	Parameter, nonnegative integer
m_d	Unit count consumed by drone sortie d	Parameter, nonnegative integer
T_r	Patrol-route hours consumed by route r	Parameter, nonnegative real (h)
T_d	Drone-sortie hours consumed by sortie d	Parameter, nonnegative real (h)
$B_{\text{patrol}}, B_{\text{drone}}, B_{\text{sensor}}$	Patrol, drone, and sensor budgets	Parameters, nonnegative integers
$H_{\text{patrol}}, H_{\text{drone}}$	Patrol-hour and drone-hour budgets	Parameters, nonnegative reals (h)

Symbol	Definition	Type
η_H, η_Q	Minimum hotspot and remote-set monitoring floors	Parameters, nonnegative reals
$Y_{g,t,k}$	Indicator that a harmful event of type k occurs in cell g at time t in the stochastic derivation	Random variable, $\{0, 1\}$
$Z_{g,t,k}$	Indicator that the harmful event of type k is successfully prevented in cell g at time t	Random variable, $\{0, 1\}$
$\ell_{g,t,k}$	Realized ecological loss if a threat of type k succeeds in cell g at time t	Derived quantity, nonnegative real
$p_{0,t,k}$	Scenario-wide incident scale factor for threat type k at time t in the stochastic derivation	Parameter, nonnegative real
L_g^0	Baseline expected loss in cell g before deployment	Derived quantity, nonnegative real
$L_{g,\omega}^0$	Baseline expected loss in cell g under scenario ω before deployment	Derived quantity, nonnegative real
L_g	Residual expected loss in cell g after deployment	Derived quantity, nonnegative real
$L_{g,\omega}$	Residual expected loss in cell g under scenario ω after deployment	Derived quantity, nonnegative real
PI	Park-wide protection index in generic form	Derived metric, $[0, 1]$
PI_ω	Park-wide protection index under scenario ω	Derived metric, $[0, 1]$
ν, α, β	Memory decay, incident-pressure amplification, and monitoring suppression parameters in the dynamic update	Parameters, nonnegative reals
$I_{g,t}$	Incident pressure at cell g and time t in the dynamic update layer	State/parameter, nonnegative real
$F_{g,t}$	Fire impact at cell g and time t in the dynamic consequence update	State/parameter, nonnegative real
ρ_F	Fire-damage rate in the dynamic consequence update	Parameter, nonnegative real
χ_g	Local ecological recovery factor in the dynamic consequence update	Parameter, $[0, 1]$
$N^\star(\tau)$	Minimum personnel needed to achieve target protection τ	Derived metric, nonnegative integer
$N_{\text{base}}, s_P, s_D, s_S, b_S$	Staffing-conversion constants for base overhead, patrol staff, drone staff, sensor support, and sensor block size	Parameters, nonnegative reals / integers

§4. Appendix D1. Derivation of the Tractable Optimization Surrogate

This appendix formalizes how the implemented MILP is obtained from a more detailed stochastic protection picture.

§4.1 Stochastic origin of the loss model

Define $Y_{g,t,k} \in \{0,1\}$ indicate that a harmful event of type k occurs in cell g during period t. Let the loss suffered if that event succeeds be

$$\ell_{g,t,k} = A_g C_{g,t,k},$$

where A_g is cell area and $C_{g,t,k}$ is the scenario-specific ecological consequence. Additionally, let $Z_{g,t,k} \in \{0,1\}$ indicate that the incident is successfully prevented. Then the realized residual loss is:

$$\widetilde{L}_{g,t,k} = \ell_{g,t,k} Y_{g,t,k}(1 - Z_{g,t,k}).$$

Taking expected values of both sides gives

$$\mathbb{E}[\widetilde{L}_{g,t,k}] = \ell_{g,t,k} \Pr(Y_{g,t,k} = 1)\Big(1 - \Pr(Z_{g,t,k} = 1 \mid Y_{g,t,k} = 1)\Big).$$

Write the incident probability as

$$\Pr(Y_{g,t,k} = 1) = p_{0,t,k}\lambda_{g,t,k},$$

where $p_{0,t,k} > 0$ is a scenario-wide scale factor and $\lambda_{g,t,k} \in [0,1]$ is the normalized spatial threat intensity. If detection probability is $D_{g,t,k}$ and conditional response success after detection is $R_{g,t,k}$, then

$$\Pr(Z_{g,t,k} = 1 \mid Y_{g,t,k} = 1) = D_{g,t,k}R_{g,t,k},$$

so

$$\mathbb{E}[\widetilde{L}_{g,t,k}] = p_{0,t,k} A_g C_{g,t,k} \lambda_{g,t,k}(1 - D_{g,t,k}R_{g,t,k}).$$

Proposition (1)

Multiplying every cell loss in a fixed scenario by the same positive constant does not change the optimal asset allocation.

Proof. If every feasible objective value is multiplied by the same positive constant $p_{0,t,k}$, the ordering of feasible solutions is unchanged! Therefore the optimizer selects the same decision vector. □

This is why the model can effectively drop $p_{0,t,k}$ and work with the relative baseline loss field

$$L^0_{g,t,k} = A_g C_{g,t,k}\lambda_{g,t,k}.$$

So L^0 is best interpreted as a *relative expected-loss surface*, not as a literal count of incidents.

§4.2 Product $A_g C_g \lambda_g$ and why it is the right baseline object

Each factor within has a distinct role.

- A_g scales loss by how much physical territory the cell represents after clipping.
- C_g measures how ecologically costly a successful incident would be in that cell.
- λ_g measures how likely that cell is to experience the relevant threat.

The product is therefore the simplest quantity that increases when either ecological importance or incident likelihood increases. A sum such as $C_g + \lambda_g$ would fail this test: it would treat a high-value safe cell and a low-value dangerous cell as interchangeable. The product, however, does not.

§4.3 Exact detection and nonlinearity

If patrol, drone, and sensor detections are conditionally independent given an incident (i.e., the success of a detection from a drone is completely independent of patrol), the natural exact detection expression is

$$D^{\text{exact}}_{g,t,k} = 1 - \left(1 - P^{\text{patrol}}_{g,t,k}\right)\left(1 - P^{\text{drone}}_{g,t,k}\right)\left(1 - P^{\text{sensor}}_{g,t,k}\right).$$

This is by evaluating the complementary probability in which none (0) detections occur.

Expanding gives

$$\begin{aligned} D^{\text{exact}}_{g,t,k} = {} & P^{\text{patrol}}_{g,t,k} + P^{\text{drone}}_{g,t,k} + P^{\text{sensor}}_{g,t,k} \\ & - P^{\text{patrol}}_{g,t,k} P^{\text{drone}}_{g,t,k} - P^{\text{patrol}}_{g,t,k} P^{\text{sensor}}_{g,t,k} - P^{\text{drone}}_{g,t,k} P^{\text{sensor}}_{g,t,k} \\ & + P^{\text{patrol}}_{g,t,k} P^{\text{drone}}_{g,t,k} P^{\text{sensor}}_{g,t,k}. \end{aligned}$$

Evaluating, we recognize that the additive first-order terms reward extra monitoring, whereas the quadratic and cubic terms encode overlap and saturation. Those interaction terms are precisely what turn the selection problem into a mixed-integer nonlinear program once asset-choice variables are introduced.

For small to moderate component probabilities, one has the local approximation

$$D^{\text{exact}}_{g,t,k} = P^{\text{patrol}}_{g,t,k} + P^{\text{drone}}_{g,t,k} + P^{\text{sensor}}_{g,t,k} + O\left(P^{\text{patrol}}_{g,t,k} P^{\text{drone}}_{g,t,k} + P^{\text{patrol}}_{g,t,k} P^{\text{sensor}}_{g,t,k} + P^{\text{drone}}_{g,t,k} P^{\text{sensor}}_{g,t,k}\right).$$

Because it maintains the first-order structure of the precise detection expression, an additive monitoring score makes sense, mathematically.

§4.4 Coverage kernels used in the implementation

As discussed in the main paper, the three asset classes do not share the same geometry, so they are assigned different influence kernels.

Patrol routes. For route polyline Γ_r,

$$\phi_{rg} = \alpha_P \exp\left(-\frac{\text{dist}(c_g, \Gamma_r)}{\rho_P}\right) \cdot \mathbf{1}\left\{\alpha_P \exp\left(-\frac{\text{dist}(c_g, \Gamma_r)}{\rho_P}\right) \geq \varepsilon_P\right\}.$$

This makes sense, and also makes routes corridor assets: strongest on-route, weaker nearby, and negligible far away.

Drone sorties. For anchor point a_d and coverage radius R_d,

$$\psi_{dg} = \alpha_D \max\left\{0, 1 - \frac{\|c_g - a_d\|}{R_d}\right\} \cdot \mathbf{1}\left\{\alpha_D \max\left(0, 1 - \frac{\|c_g - a_d\|}{R_d}\right) \geq \varepsilon_D\right\}.$$

This yields a quite broad but shallow footprint.

Sensor sites. If $\text{vis}_{sg} \in \{0, 1\}$ is the DEM-derived visibility indicator,

$$\chi_{sg} = \text{vis}_{sg} \max\left\{\chi_{\min}, \chi_0 \exp\left(-\frac{\text{dist}(s, c_g)}{\rho_S}\right)\right\} \mathbf{1}\{\text{dist}(s, c_g) \leq R_S\}.$$

As wished, the sensor kernel is persistent, yet it is limited by line-of-sight and radius.

All three kernels are nonnegative, sparse after truncation, and directly precomputable, making the later MILP possible.

§4.5 Response-time transformation

Best-response time is computed from the mobility graph and terminal off-road spur:

$$\tau_g = \min_b \tau_{bg}.$$

Response success is then modeled by

$$R_g = \frac{1}{1 + \exp(\kappa(\tau_g - \tau_0))}.$$

This choice has three desirable properties:

- $R_g \in (0, 1)$ automatically;
- $R(\tau_0) = 1/2$, so τ_0 is an interpretable operational breakpoint;
- $R'(\tau) = -\kappa R(\tau)(1 - R(\tau)) < 0$, so slower response always hurts.

Proposition (2)

The logistic response curve is strictly decreasing in τ and has maximum slope magnitude at $\tau = \tau_0$.

Proof. This is a property of logistic functions. We may see by differentiating:

$$\frac{dR}{d\tau} = -\kappa R(\tau)(1 - R(\tau)).$$

Since $0 < R(\tau) < 1$ and $\kappa > 0$, the derivative is always negative. The expression $R(1 - R)$ is maximized at $R = 1/2$ by the AM-GM equality case, which occurs exactly at $\tau = \tau_0$. □

§4.6 Scenario-static simplification and monitoring aggregation

The conceptual model is indexed by time t and threat type k. For a strategic planning paper, solving the fully time-indexed problem would add large state and scheduling complexity. We therefore collapse the dynamic structure into scenario-specific baseline loss fields. For each scenario ω,

$$L^0_{g,\omega} = A_g C_{g,\omega} \lambda_{g,\omega}.$$

The strategic model then solves one optimization per scenario.

Next, instead of preserving separate patrol, drone, and sensor detection probabilities inside the optimizer, we aggregate them into one monitoring score:

$$M_g = \sum_r \phi_{rg} y_r + \sum_d \psi_{dg} z_d + \sum_s \chi_{sg} x_s.$$

This is the implemented first-order monitoring quantity. It is linear in the binary decisions and therefore solver-friendly.

§4.7 The detection surrogate

Let $0 = b_0 < b_1 < \cdots < b_L$ be chosen monitoring levels, and $0 = v_0 < v_1 < \cdots < v_L < 1$ be the corresponding detection values. These level-value pairs define the shape of the surrogate curve.

For each interval $[b_\ell, b_{\ell+1}]$, define the slope

$$m_\ell = \frac{v_{\ell+1} - v_\ell}{b_{\ell+1} - b_\ell},$$

and require $m_0 \geq m_1 \geq \cdots \geq m_{L-1} \geq 0$. This ensures that the surrogate is increasing but exhibits diminishing returns, since the slopes of the secant lines decrease over time.

Next, for each segment define the affine function (the equation for the secant line)

$$u_\ell(M) = v_\ell + m_\ell(M - b_\ell).$$

We then impose the linear constraints

$$S_g \leq u_\ell(M_g) \qquad \forall \ell,$$

together with

$$0 \leq S_g \leq v_L.$$

Because the objective is to maximize S_g, the optimizer will set

$$S_g = \min_\ell u_\ell(M_g),$$

so S_g becomes the intended concave piecewise-linear surrogate for detection.

In the present calibration, the S_g is capped at 0.45. This reflects the fact that S_g is only a partial detection proxy (not a full success probability.) Avoided loss still depends on the response factor R_g (consult section 6.1), so high monitoring alone does not imply full protection.

§4.8 Why this works

Our provided construction is rigorous for three main reasons.

1. It is **monotonically true**: more monitoring never lowers S_g.
2. It has **diminishing returns**: decreasing secant slopes show that with increased monitoring, detection gradually levels off
3. It is **correct in probability limit**: the strict upper bound $v_L < 1$ means that the model never assumes detection is certain

We emphasize that this piecewise-linear approximation makes the model mathematically cleaner, keeps it fully within the MILP framework, and preserves the two main behaviors the optimization needs – namely, (1) protection increases as monitoring increases, (2) but extra monitoring provides smaller gains over time.

Proposition (3)

Assume that all assets contribute a nonnegative amount to each cell, meaning $\phi_{rg}, \psi_{dg}, \chi_{sg} \geq 0$, and that the piecewise-linear detection curve u_ℓ is comprised of nonincreasing secant slopes. Then for any fixed cell g, the avoided-loss contribution

$$f_g = L_g^0 R_g S_g$$

satisfies two properties:

1. Increasing the monitoring contribution of any selected asset cannot decrease avoided loss.
2. As the total monitoring score M_g grows, each additional increase in monitoring yields a smaller gain in avoided loss than the previous one.

§4.9 Concave detection envelope

Choose anchor pairs (b_ℓ, v_ℓ) with

$$0 = b_0 < b_1 < \cdots < b_L, \qquad 0 = v_0 < v_1 < \cdots < v_L < 1,$$

and decreasing secant slopes

$$m_\ell = \frac{v_{\ell+1} - v_\ell}{b_{\ell+1} - b_\ell}, \qquad m_0 \geq m_1 \geq \cdots \geq m_{L-1} \geq 0.$$

Define the affine lines

$$u_\ell(M) = v_\ell + m_\ell(M - b_\ell).$$

The surrogate variable S_g is constrained by

$$0 \leq S_g \leq v_L, \qquad S_g \leq u_\ell(M_g) \quad \forall \ell.$$

Proposition (4)

Suppose the slopes $\{m_\ell\}$ are nonincreasing and the objective coefficient of S_g is nonnegative. Then at optimality,

$$S_g = \min_\ell u_\ell(M_g),$$

which is a monotone, concave, bounded function of M_g passing through the anchor values.

Proof. The feasible set for (M_g, S_g) is the intersection of half-planes below the lines u_ℓ. Since S_g enters the objective with a nonnegative coefficient, the optimizer pushes S_g upward as far as feasibility allows. Therefore the largest feasible value is attained, namely the lower envelope of the affine bounds. Because the segment slopes are nonincreasing, that lower envelope is concave. The bound $S_g \leq v_L < 1$ keeps the surrogate probability-safe (3 in Subsection 4.8). □

In the current calibration the breakpoints are

$$(0, 0),\ (1, 0.12),\ (2, 0.22),\ (4, 0.34),\ (6, 0.45),$$

so the model explicitly encodes diminishing returns and caps detection below perfect certainty.

§4.10 Resulting MILP

Once $L^0_{g,\omega}$, R_g, ϕ_{rg}, ψ_{dg}, and χ_{sg} are all precomputed, the allocation problem becomes

$$\max \sum_{g\in G} L^0_{g,\omega} R_g S_g$$

subject to

$$M_g = \sum_r \phi_{rg} y_r + \sum_d \psi_{dg} z_d + \sum_s \chi_{sg} x_s,$$

the envelope inequalities for S_g, and the budget/hour constraints.

To control spatial equity and baseline coverage, the model also supports two regional floor constraints. The first is a hotspot-monitoring condition that prevents the optimizer from ignoring top-decile loss cells. If $H = \left\{g \in G : \ L^0_g \geq Q_{0.9}(L^0)\right\}$, the implemented hotspot floor is

$$\sum_{g\in H} M_g \geq \eta_H |H|,$$

with $\eta_H = 0.20$ in the present baseline. The second is an optional fairness condition for a designated remote-cell set $Q \subseteq G$ chosen from the upper tail of response times. If activated, the fairness floor is

$$\sum_{g\in Q} M_g \geq \eta_Q |Q|.$$

Proposition (5)

The implemented allocation problem is a MILP.

Proof. The objective is linear in S_g because the coefficients $L^0_{g,\omega} R_g$ are precomputed constants. The monitoring equations are linear equalities, the envelope restrictions are linear inequalities, and the budget constraints are linear. The only integrality conditions are the binary asset decisions. Hence the model is a mixed-integer linear program. □

§4.11 Staffing interpretation

The staffing analysis in the paper is also derived rather than ad hoc. If B_P is the patrol-route budget, then the personnel conversion used by the implementation is

$$N_{\text{staff}}(B_P) = N_{\text{base}} + s_P B_P + s_D B_D + s_S \left\lceil \frac{B_S}{b_S} \right\rceil,$$

where:

- N_{base} is standing base overhead,
- s_P is personnel per patrol route,
- s_D is personnel per drone unit,
- s_S is support staff per sensor block,
- b_S is the sensor block size.

The minimum staffing needed to reach target protection τ is then

$$N^\star(\tau) = \min\{N_{\text{staff}}(B_P) : PI^\star(B_P) \geq \tau\},$$

where $PI^{\star}(B_P)$ is the optimal protection index under that patrol budget and the fixed baseline drone/sensor assumptions. For this reason, the main paper's staffing curve is more than solely descriptive. It is a picture of a systematic increase in patrol capacity combined with repeated optimization.

> **Remark** (1)**.** We think that the next mathematical advancements would be a fully time-indexed deployment model, a more accurate SOS2 approximation of the detection curve, or a McCormick-relaxed expansion of the exact multiplicative detection formula if greater computational complexity were both feasible and acceptable. These are improvements in accuracy, but they are not necessary for a strategic model that can be defended.

§5. Appendix D2. Other Derivations

This subsection complements Appendix D1. This is by collecting derivations that support the ecological-consequence layer, threat layer construction, combined scenario aggregation, dynamic update equations, and remote-sector fairness metrics. These details are very useful for underlying mathematical compeleteness.

§5.1 Seasonal consequence construction

For each cell g and season $\sigma \in \{\text{dry}, \text{wet}\}$, define the raw species-use score

$$\widetilde{C}_{g,\sigma} = \omega_r q^{(r)}_{g,\sigma} + \omega_e q^{(e)}_{g,\sigma} + \omega_h q^{(h)}_{g,\sigma}, \qquad \omega_r > \omega_e > \omega_h > 0, \tag{1}$$

where the weights reflect the priority ordering (relative) of rhino, elephant, and general herbivore protection in the setting of anti-poaching, respectively.

The implemented species-use fields are

$$q^{(r)}_{g,\sigma} = \exp\left(-\frac{d_{g,W}}{\rho^{(r)}_{W,\sigma}}\right) \exp\left(-\frac{d_{g,\text{fence}}}{\rho_F}\right), \tag{2}$$

$$q^{(e)}_{g,\sigma} = \exp\left(-\frac{d_{g,W}}{\rho^{(e)}_{W,\sigma}}\right), \tag{3}$$

$$q^{(h)}_{g,\sigma} = \exp\left(-\frac{d_{g,W}}{\rho^{(h)}_{W,\sigma}}\right), \tag{4}$$

where $d_{g,W}$ represents the distance from cell g to the nearest waterhole and $d_{g,\text{fence}}$ is the distance to the park boundary.

¶ Why dry-season value is more concentrated. Suppose $\rho^{(u)}_{W,\text{dry}} < \rho^{(u)}_{W,\text{wet}}$ for a species group $u \in \{r, e, h\}$. Then for every $d > 0$,

$$\exp\left(-\frac{d}{\rho^{(u)}_{W,\text{dry}}}\right) < \exp\left(-\frac{d}{\rho^{(u)}_{W,\text{wet}}}\right). \tag{5}$$

Hence the dry-season kernel decays more quickly with distance from water, which concentrates more of the consequence mass near the waterhole network.

¶ Normalization. For any nonconstant raw field $\widetilde{X}_g$, define

$$\text{norm}(\widetilde{X}_g) = \frac{\widetilde{X}_g - \min_{h\in G} \widetilde{X}_h}{\max_{h\in G} \widetilde{X}_h - \min_{h\in G} \widetilde{X}_h}. \tag{6}$$

Accordingly,

$$C_{g,\sigma} = \text{norm}(\widetilde{C}_{g,\sigma}). \tag{7}$$

This places all consequence layers on a common $[0, 1]$ scale before they are multiplied by threat intensity.

¶ Wildfire consequence. The fire-consequence field is constructed separately:

$$\widetilde{C}_{g,\text{fire}} = B_g\left(0.35 + 0.65\exp\left(-\frac{d_{g,W}}{\rho_{W,F}}\right)\right), \tag{8}$$

$$C_{g,\text{fire}} = \text{norm}(\widetilde{C}_{g,\text{fire}}). \tag{9}$$

The constant term 0.35 prevents the field from degenerating to zero far from water, while the exponential term still allows water-adjacent biomass zones to carry higher ecological value. Again, logical in application.

§5.2 Threat-kernel construction and monotonicity

The dry-season poaching, wet-season poaching, intrusion, and wildfire layers are first built as raw nonnegative scores and then normalized.

For poaching,

$$\widetilde{\lambda}_{g,\text{dry}} = \beta_r q^{(r)}_{g,\text{dry}} + \beta_G\exp\left(-\frac{d_{g,\text{gate}}}{\rho_G}\right) + \beta_R\exp\left(-\frac{d_{g,\text{road}}}{\rho_R}\right) + \beta_F\exp\left(-\frac{d_{g,\text{fence}}}{\rho_F}\right), \tag{10}$$

$$\widetilde{\lambda}_{g,\text{wet}} = \beta_r q^{(r)}_{g,\text{wet}} + \beta_G\exp\left(-\frac{d_{g,\text{gate}}}{\rho_G}\right) + \beta_R\exp\left(-\frac{d_{g,\text{road}}}{\rho_R}\right) + \beta_F\exp\left(-\frac{d_{g,\text{fence}}}{\rho_F}\right). \tag{11}$$

For intrusion,

$$\widetilde{\lambda}_{g,\text{intr}} = \eta_G\exp\left(-\frac{d_{g,\text{gate}}}{\rho_{IG}}\right) + \eta_F\exp\left(-\frac{d_{g,\text{fence}}}{\rho_{IF}}\right) + \eta_R\exp\left(-\frac{d_{g,\text{road}}}{\rho_{IR}}\right). \tag{12}$$

For wildfire,

$$\begin{aligned}\widetilde{\lambda}_{g,\text{fire}} = {} & \zeta_1 B_g + \zeta_2\left(1-\exp\left(-\frac{d_{g,\text{ctr}}}{\rho_{\text{ctr},F}}\right)\right) + \zeta_3\exp\left(-\frac{d_{g,\text{fence}}}{\rho_{\text{fence},F}}\right) + \zeta_4 H_g + \zeta_5 J_g \\ & + \zeta_6 X_g + \zeta_7 D_t(0.5+0.5B_g) + \zeta_8 W_t X_g + \zeta_9 S_g^{\text{slope}}.\end{aligned} \tag{13}$$

The final hazard layers are then

$$\lambda_{g,\omega} = \text{norm}(\widetilde{\lambda}_{g,\omega}) \qquad \text{for } \omega \in \{\text{dry}, \text{wet}, \text{intr}, \text{fire}\}. \tag{14}$$

¶ Why the distance kernels behave correctly. For any decay length $\rho > 0$,

$$\frac{\partial}{\partial d}\exp\left(-\frac{d}{\rho}\right) = -\frac{1}{\rho}\exp\left(-\frac{d}{\rho}\right) < 0. \tag{15}$$

Thus gate-, road-, and fence-proximity terms decrease with distance, so closer access points raise threat as intended. Likewise,

$$\frac{\partial}{\partial d}\left(1-\exp\left(-\frac{d}{\rho}\right)\right) = \frac{1}{\rho}\exp\left(-\frac{d}{\rho}\right) > 0, \tag{16}$$

so the center-distance term in the fire layer increases away from the center surrogate.

§5.3 Combined weighted scenario derivation

Let $J = \{\text{dry}, \text{wet}, \text{intr}, \text{fire}\}$ index the components of the combined scenario, and let $w_j \geq 0$ be the weight attached to component $j \in J$. For each component define

$$L^0_{g,j} = A_g C_{g,j} \lambda_{g,j},$$

where, in the present Etosha calibration, $C_{g,\text{intr}}$ is taken to be the dry-season consequence field.

The combined baseline loss in cell g is

$$L^0_{g,\text{comb}} = \sum_{j \in J} w_j L^0_{g,j}.$$

If $R_{g,j}$ denotes the response coefficient used for component j, the combined avoided loss is

$$\text{AvoidedLoss}_{g,\text{comb}} = \sum_{j \in J} w_j L^0_{g,j} R_{g,j} S_g,$$

and the residual loss is

$$L_{g,\text{comb}} = L^0_{g,\text{comb}} - \text{AvoidedLoss}_{g,\text{comb}}.$$

Summing over all cells gives

$$L^0_{\text{comb}} = \sum_{g \in G} L^0_{g,\text{comb}}, \qquad L_{\text{comb}} = \sum_{g \in G} L_{g,\text{comb}},$$

hence

$$PI_{\text{comb}} = \frac{\sum_{g \in G} \sum_{j \in J} w_j L^0_{g,j} R_{g,j} S_g}{\sum_{g \in G} \sum_{j \in J} w_j L^0_{g,j}}.$$

If a particular calibration uses a common response coefficient across all components, then one simply sets $R_{g,j} \equiv R_g$.

§5.4 Dynamic update layer

The static optimization solves for one deployment under one calibrated scenario. To represent continued pressure over time, the model adds a deterministic update layer.

For hazard intensity, let $\Pi_{[0,1.5]}(x) = \min\{1.5, \max\{0, x\}\}$ denote clipping to the implemented state range. Let

$$\widehat{I}_{g,t} = s_I \frac{\text{residual\_loss}_{g,t}}{\max_{h \in G} \text{residual\_loss}_{h,t}}, \qquad \widehat{M}_{g,t} = \frac{M_{g,t}}{\max_{h \in G} M_{h,t}},$$

where s_I is the incident-pressure scaling constant. The implemented hazard update is

$$\lambda_{g,t+1} = \Pi_{[0,1.5]}\Big((1-\nu)\lambda_{g,t} + \alpha \widehat{I}_{g,t} - \beta \widehat{M}_{g,t}\Big).$$

Thus unresolved residual loss raises future hazard, while sustained monitoring suppresses it.

For the combined scenario, ecological consequence is updated by

$$C_{g,t+1} = \Pi_{[0,1.5]}\Big(C_{g,t}(1 - \rho_F H^{\text{fire}}_g) + \gamma(1 - C_{g,t})\Big),$$

where $H^{\text{fire}}_g \in [0,1]$ is the fixed normalized fire-hazard field, ρ_F is the fire-damage rate, and γ is the global recovery rate. The first term reduces current consequence after fire pressure, while the second allows gradual recovery toward the undamaged state.

§5.5 Remote-sector fairness metrics

Let $Q \subseteq G$ denote the designated remote-cell set. A natural remote protection index, following, is:

$$PI_Q = \frac{\sum_{g \in Q} L_g^0 S_g R_g}{\sum_{g \in Q} L_g^0}. \tag{17}$$

Realize that this is the same fraction of avoided loss as the park-wide protection index, but restricted to the remote subset.

Two additional summary indicators are

$$\overline{M}_Q = \frac{1}{|Q|} \sum_{g \in Q} M_g, \tag{18}$$

$$\overline{\tau}_Q = \frac{1}{|Q|} \sum_{g \in Q} \tau_g. \tag{19}$$

These quantities separate three distinct fairness questions:

- whether remote cells are being watched ($\overline{M}_Q$),
- whether they are actually reachable in time ($\overline{\tau}_Q$),
- whether they are actually protected in avoided-loss terms (PI_Q).

§6. Appendix E. Calibration and Parameter Grounding

Table 9: Calibration and Parameter Grounding

Item	Current value	Evidence or rationale
Grid size	5.000 km	Resolves hotspot belts while keeping repeated optimization and stress tests tractable.
Dry waterhole decay	Not fixed, but shorter than wet	Reflects stronger dry-season concentration around permanent water on the official park map and field narratives (*Map of Etosha*, n.d.).
Response midpoint τ_0	2.000 h	Operational threshold chosen to distinguish actionable from weak interception windows; sensitivity tests in main paper show conclusions are stable under midpoint shifts.
Road speeds	30–45 km/h by class	Conservative park-operational assumption, set below public-road intuition to account for observation stops and wildlife caution.
Off-road speed	Tobler function	Standard and famous terrain-dependent walking model (Tobler, 1993).
Terrain source	OpenTopography DEM	Globally accessible elevation source suitable for transfer (OpenTopography, 2021).
Fire weather	NASA POWER	Supplies precipitation, humidity, temperature, and wind summaries (NASA, 2025).
Fire detections	NASA FIRMS	Supplies recent and historical active-fire observations (NASA, n.d.-a).
Detection cap	0.45	Conservative upper bound ensuring $S_g < 1$ before multiplication by response.
Detection anchors	$(0, 0), (1, 0.12)$ $(2, 0.22), (4, 0.34)$ $(6, 0.45)$	Piecewise-linear diminishing-returns envelope for converting monitoring intensity into detection probability.
Hotspot quantile	0.90	Identifies the top decile of cells by baseline loss as managerial hotspots.
Patrol team size	6	Operational planning assumption used to translate route capacity into headcount.
Drone team size	2	Pilot-plus-support-operator assumption.
Sensor support block	1 per 4 sensors	Maintenance and oversight overhead assumption.
Base staff overhead	12	Standing operational overhead included in staffing conversion.
Fairness switch	Off in baseline	Lets the paper quantify the pure hotspot-first optimum before imposing a service-floor policy.

Table 10: Selected baseline parameters for the Etosha analysis

Quantity	Value	Interpretation
Grid size	5.000 km	resolves waterhole structure while keeping repeated solves tractable
Patrol budget	10	baseline route selections
Drone budget	4	baseline drone sorties
Sensor budget	12	baseline persistent sites
Patrol hours budget	80.000 h	aggregate route-hours constraint
Drone hours budget	24.000 h	aggregate sortie-hours constraint
Response midpoint τ_0	2.000 h	half-success response threshold
Response steepness κ	2.2	controls how sharply interception decays with time
Combined weights	0.40/0.35/ 0.15/0.10	dry/wet/intrusion/fire scenario mix
Hotspot quantile	0.90	top-decile loss cells used as hotspot set
Total staff ceiling	295	park-wide ceiling from the IMMC brief

Calibration maturity

We believe the poaching and response modules are the most strongly grounded parts of the model. The wildfire module is improved by real fire-weather and detection inputs, but its calibration is less mature because it still lacks dedicated fuel-load and burn-scar layers.

§7. Appendix F. Full Results Tables

Table 11: Scenario results

Scenario	Base loss	Residual loss	Avoided loss	PI	Hotspot PI
Dry poaching	3412.796	2906.767	506.028	0.148	0.244
Wet poaching	4195.320	3624.943	570.377	0.136	0.253
Combined weighted	4127.833	3593.394	534.439	0.129	0.250

Table 12: Response-time summary from the best-base service surface

Metric	Value
Median best response time	1.935 h
90th percentile	4.093 h
95th percentile	5.437 h
99th percentile	8.154 h
Cells within 4 h	89.6%
Cells within 6 h	96.6%
Unreachable outlier cells	3

Table 13: Selected staffing requirements

Scenario	PI = 0.12	PI = 0.14	PI = 0.16
Dry poaching	71	83	95
Wet poaching	71	95	107
Combined weighted	83	95	119

Table 14: Deterministic sensitivity results for the combined scenario

Case	PI	Absolute change	Relative change
Baseline	0.129472	0.000000	0.0%
Poaching surge	0.131274	+0.001802	+1.39%
Fire year	0.122425	-0.007047	-5.44%
Staff shortfall	0.108977	-0.020495	-15.83%
Drone loss	0.124508	-0.004964	-3.83%
Sensor dropout	0.126824	-0.002648	-2.05%
Slow response	0.118032	-0.011440	-8.84%

Why the sensitivity ranking behaves this way. The results in Table 14 are indeed logical. The avoided-loss formula is

$$\text{AvoidedLoss}_g = L_g^0 R_g S_g.$$

Patrols, drones, and sensors raise S_g, but if slow response pushes R_g downward, the marginal value of those assets collapses. This is precisely why slow response and staff shortfall are more damaging than moderate sensor or drone loss.

Table 15: Monte Carlo summary statistics

Scenario	Mean PI	Std. dev.	Variance	5th pct.	95th pct.	Pr(PI < baseline)
Dry poaching	0.146865	0.005623	0.000032	0.138126	0.155267	0.625
Wet poaching	0.134629	0.005274	0.000028	0.126416	0.142593	0.625
Combined weighted	0.127652	0.005127	0.000026	0.119783	0.134355	0.667

Table 16: Fairness metrics

Scenario	Overall PI	Remote PI	Overall mean response	Remote mean response
Dry poaching	0.148274	3.14×10^{-8}	5.227 h	34.633 h
Wet poaching	0.135956	2.58×10^{-8}	5.227 h	34.633 h
Combined weighted	0.129472	7.66×10^{-8}	5.227 h	34.633 h

Table 17: Dynamic scenario (combined) trajectory

Period	Base loss	Avoided loss	Protection index
0	3913.003	527.631	0.134840
1	3646.192	464.966	0.127521
2	3405.494	408.356	0.119911
3	3185.638	356.431	0.111887
4	2982.021	308.096	0.103318
5	2779.500	261.321	0.094017
6	2580.555	216.633	0.083948
7	2394.412	176.926	0.073891
8	2225.020	143.145	0.064334

Operational placement for combined To make the deployment actionable (i.e., rather than purely visual), Tables 18–20 show the selected combined scenario patrol staging bases, drone anchor cells, and sensor locations. The patrol coordinates indicate the location from which a moving patrol route is initiated; the geometry of the route in the main text shows the route itself, and the hours indicate the approximate duration of a single pass along this route

Table 18: Combined-scenario selected patrol routes (EPSG:32733). Coordinates mark staging bases; hours are per continuous pass of the displayed route geometry.

Base and patrol type	Easting (m)	Northing (m)	Waypoint waterholes	Hours/pass
CAMP001, base-cluster patrol	444644	7897999	WH002 – WH001	3.09
CAMP001, base-cluster patrol	444644	7897999	WH003 – WH006	3.45
CAMP002, base-cluster patrol	485359	7902321	WH009 – WH010	4.52
CAMP002, base-cluster patrol	485359	7902321	WH008 – WH011	3.93
CAMP003, base-cluster patrol	596595	7879094	WH031 – WH033	2.64
CAMP004, base-cluster patrol	654815	7894630	WH033 – WH031	4.13
CAMP004, base-cluster patrol	654815	7894630	WH032 – WH030	3.72
CAMP006, base-cluster patrol	681929	7936051	WH051 – WH052	4.40
CAMP006, base-cluster patrol	681929	7936051	WH050 – WH049	3.27
GATE003, entry-screen patrol	684131	7953511	WH055 – WH056	3.44

Table 19: Combined-scenario selected drone sortie anchors (EPSG:32733)

Anchor cell	Easting (m)	Northing (m)	Radius (km)	Hours
Cell 561	619445	7871906	16	2.0
Cell 503	604445	7876906	16	2.0
Cell 71	464445	7881906	16	2.0
Cell 911	699445	7926906	16	2.0

Table 20: Combined-scenario selected sensor sites (EPSG:32733). Strategic sites are optimizer-generated fixed candidates not tied to named camps or waterholes.

Site	Type	Easting (m)	Northing (m)
CAMP002	camp	485359	7902321
CAMP003	camp	596595	7879094
CAMP005	camp	704428	7919276
WH028	waterhole	598842	7864022
WH049	waterhole	703301	7918101
WH031	waterhole	618867	7872354
WH027	waterhole	603377	7876241
WH008	waterhole	458994	7904489
WH010	waterhole	476870	7882207
Strategic 022	strategic	599717	7862926
Strategic 023	strategic	713529	7922066
Strategic 025	strategic	444445	7901379

§8. Appendix G. Transfer Workflow and Park Adaptation Details

A transfer checklist

Keep fixed: the technicals—loss-based objective, hybrid grid-graph structure, terrain-aware response logic, multi-asset monitoring, and staffing/robustness workflow.
Regenerate automatically where possible: boundary, roads, elevation, slope, line of sight, fire detections, and fire-weather context from public data.
Re-estimate locally: ecological attractors, threat weights, operating bases, seasonal closures, flood or snow constraints, budget limits, and fairness policy.

1. Park definition
Boundary polygon, local management scope, staffing/budget assumptions. The calibrative values and macroscopic park-specific features should also be adjusted/added.

↓

2. Automated public-data ingestion
OpenStreetMap roads and tagged nodes; OpenTopography DEM; NASA FIRMS fire detections; NASA POWER weather summaries

↓

3. Regenerated maps
Grid, terrain-aware mobility graph, slope and visibility layers, response-time matrix, baseline fire context

↓

4. Local calibration
Species or habitat value layers, focal threat covariates, corridor/refuge nodes, closure rules, fairness targets

↓

5. Optimization and outputs
Patrol, drone, and sensor allocation; protection index; staffing curve; sensitivity and robustness analyses

Figure 14: Step-by-step guide for cross-park transfer. The boundary and public geospatial services rebuild the generic spatial backbone automatically, while ecological semantics and operating policy are recalibrated locally.

Table 21: Park-specific recalibration checklist for the two comparison parks

Park	Replace consequence with	Replace threat and mobility with	Expected deployment shift
Yellowstone	wetland and riparian habitat, bear/bison interaction zones, visitor-sensitive ecological areas (National Park Service, 2025b; National Park Service, 2025c)	wildfire-heavy threat weights, entrance and campground pressure, snow-season road closures, steep-terrain response penalties (*Fire*, 2026; *Park Roads*, 2026)	more sensors and patrol emphasis near entrances, developed nodes, canyon junctions, and high-conflict wildlife-viewing corridors
Kaziranga	rhino and megafauna habitat, flood refuges, animal-corridor value, high-ground sanctuary nodes (UNESCO World Heritage Centre, n.d.; Government of Assam, 2021)	poaching-access and highway-corridor pressure, inundation-sensitive traversal, seasonal camp disruption, refuge-oriented redeployment (Government of Assam, 2022; District Disaster Management Authority (DDMA), Golaghat, 2025)	stronger chokepoint protection, more seasonal relocation of patrol effort, and more bases or sensors near crossings, corridors, and elevated refuges

§9. Appendix H. Mathematical Extensions

Each major limitation in the main paper points to a mathematically precise future extension.

1. **Wildfire calibration upgrade.** Replace the present hybrid fire layer by a fully empirical model with fuel-load rasters, burn-scar history, drought indices, and explicit spread covariates. This requires linking more external APIs and heavy calculation.
2. **Shift-indexed optimization.** Replace the scenario-static MILP by a genuinely time-indexed deployment model with explicit team rosters and redeployment decisions. This was not feasible in the present paper because it would require a much larger state space, substantially more calibration data on shifts and movement timing, and far far greater computational effort than was realistic for this contest setting.
3. **Fairness-constrained optimization.** Activate a formal service-floor constraint or a multi-objective loss-equity trade-off and trace the resulting Pareto frontier. This is especially useful for decision makers who need to guarantee a minimum level of protection in under-served sectors of their park, such as potentially farmlands.
4. **Stochastic traversal and failures.** Replace deterministic travel and homogeneous agents with uncertainty in vehicle performance, communications, flooding, and human capacity.
5. **Adversarial learning.** Couple the protection model to a game-theoretic or adaptive adversary model in which poachers react strategically to repeated deployments. A natural next step would be to let threat intensity update from observed patrol patterns and then solve for deployments that remain effective even when adversaries learn and adapt over time.
6. **Cross-park reruns.** Carry the same architecture to additional parks with genuine recalibration. Check Appendix I, right below this section, for an empirical rerun on two other national parks.

§10. Appendix I. Rerunning the model on two national parks

This appendix goes beyond the conceptual transferring notes from the main paper to record the actual cross-park rerun. We reran the same optimization model on two other national parks – **Yellowstone National Park**, USA (North America), and **Kaziranga National Park**, India (Asia), using separate park profiles, separate raw/processed/output directories, and newly regenerated geospatial inputs.

Similarly to Etosha for both parks, the boundary, roads, and tagged point features were fetched from OpenStreetMap, then converted into park-specific node sets and local covariate layers. In Yellowstone, those transferred local layers are hydrology, visitor anchors, and winter-closure segments. In Kaziranga, they are hydrology, corridor segments, refuge points, a flood mask, and supplementary south-boundary access bases which are inferred from boundary-touching roads on the same OSM. We deliberately did *not* introduce last-resort external datasets here. The reruns therefore mirror the Etosha backbone as closely as the current environment permits. In the final appendix rerun, we also activated the same OpenTopography DEM pipeline used for Etosha, so both parks now use terrain-aware road speeds and terrain-aware final off-road approach speeds (i.e., Tobler's) rather the erroneous arbitrary flat terrain assumption.

The burden from an evidence perspective is transferability rather than full operations reporting. We therefore keep only the figures that actually strengthen the applicability claim; the following:

1. the park road-access map, because reachable geometry is the operational backbone;
2. the park-specific ecological-anchor map, because neither rerun is allowed to lean on Etosha-style waterhole logic;
3. the combined baseline loss map, because it shows the new consequence-threat geometry before optimization;
4. the best-response-time map, because it is the clearest map-level check that travel realism survives transfer; and
5. the dynamic mobility graph, because it explicitly shows seasonal edge deletion under snow or flood disruption.

We omit selected-asset maps, staffing curves, and full robustness panels for these reruns because those mostly repeat the Etosha analysis rather than adding new evidence about cross-park applicability that requires extensive coding beyond the Etosha system.

Table 22: Empirical rerun summary for the two transferred parks

Park	Road segs.	Bases	Eco. anchors	Cells	Combined PI	Local mobility layer
Yellowstone	1336	9	120	408	0.141	19 winter-closure segments extracted from the OSM road graph; a park-specific UTM projection preserves the park's road-realistic open-road and winter-closure geometry.
Kaziranga	221	4	19	28	0.405	flood-sensitive mask plus 219 corridor segments, 11 crossing nodes, and 4 south-boundary access bases inferred from the same OSM road backbone.

Table 23: Transferred layers and modeled response-time diagnostics for the reruns

Park	Transferred local layers	Response-time diagnostics
Yellowstone	478 hydrology points (ponds, wetlands, and other water-linked habitat structure), 120 visitor anchors (campground, viewpoint, attractions), 19 winter-closure segments, and an OpenTopography DEM on top of the regenerated OSM road graph and 9 access bases.	With DEM enabled, summer total response is 3.08/20.68/69.90 h (median/p95/max) and winter is 3.49/20.80/69.91 h; median road-only travel is about 0.28 h in both regimes.
Kaziranga	239 hydrology points, 219 corridor segments, 11 crossing nodes (animal crossing & movement bottlenecks), 8 refuge nodes (higher-ground or safer fallback points that animals move toward during floods), 4 operational access bases, and an OpenTopography DEM; 1 base is explicit in OSM and 3 are inferred from boundary-touching roads on the same backbone.	With DEM enabled, non-flood total response is 0.88/1.95/3.60 h and flood-sensitive response is 1.31/2.16/3.61 h; median road-only travel is about 0.21 h and 0.12 h, respectively.

The response-time heatmaps shown below clip the upper 1% of the modeled distribution for readability, but the table reports the full underlying values. We would like to clarify that these protection indices are not intended as a simple cross-park ranking. The parks differ in (1) area, (2) operating-base density, and (3) provisional ecological calibration.

What actually matters is that the same objective, response logic, and optimization machinery remain numerically stable under two sharply different geometries: a snow-season mountain road network in Yellowstone and a flood-sensitive corridor system in Kaziranga.

Yellowstone rerun. Yellowstone is the clearest test that the framework is not Etosha-specific. The rerun uses 9 OSM-derived gate or access bases, 120 curated ecological anchors (i.e., park specific points the model uses to represent where ecological value or concentration is organized) built from campgrounds, viewpoints, and attractions, 19 winter-closure segments identified directly from conditional road tags and named closed-road segments, and an OpenTopography DEM for terrain-aware speed reductions. The resulting baseline loss no longer collapses toward point waterholes; instead, it organizes around the entrance-accessible corridor system, interior visitor nodes, and hydrology-rich habitat structure. The response-time and winter edge-deletion maps separate two distinct ideas cleanly: the open-road response surface shows the terrain-aware total reach of the current access backbone, while the winter graph deletes seasonally unavailable edges before the response matrix is recomputed. Quantitatively, the summer regime has median/p95 total response of about 3.08 and 20.68 hours, respectively, while the winter regime shifts those to about 3.49 and 20.80 hours, respectively. Median road-only travel remains about 0.28 hours in both regimes, so the heavy tail comes from long, rugged final approaches into Yellowstone's road-sparse backcountry once terrain is charged explicitly. Please consult the following figures.

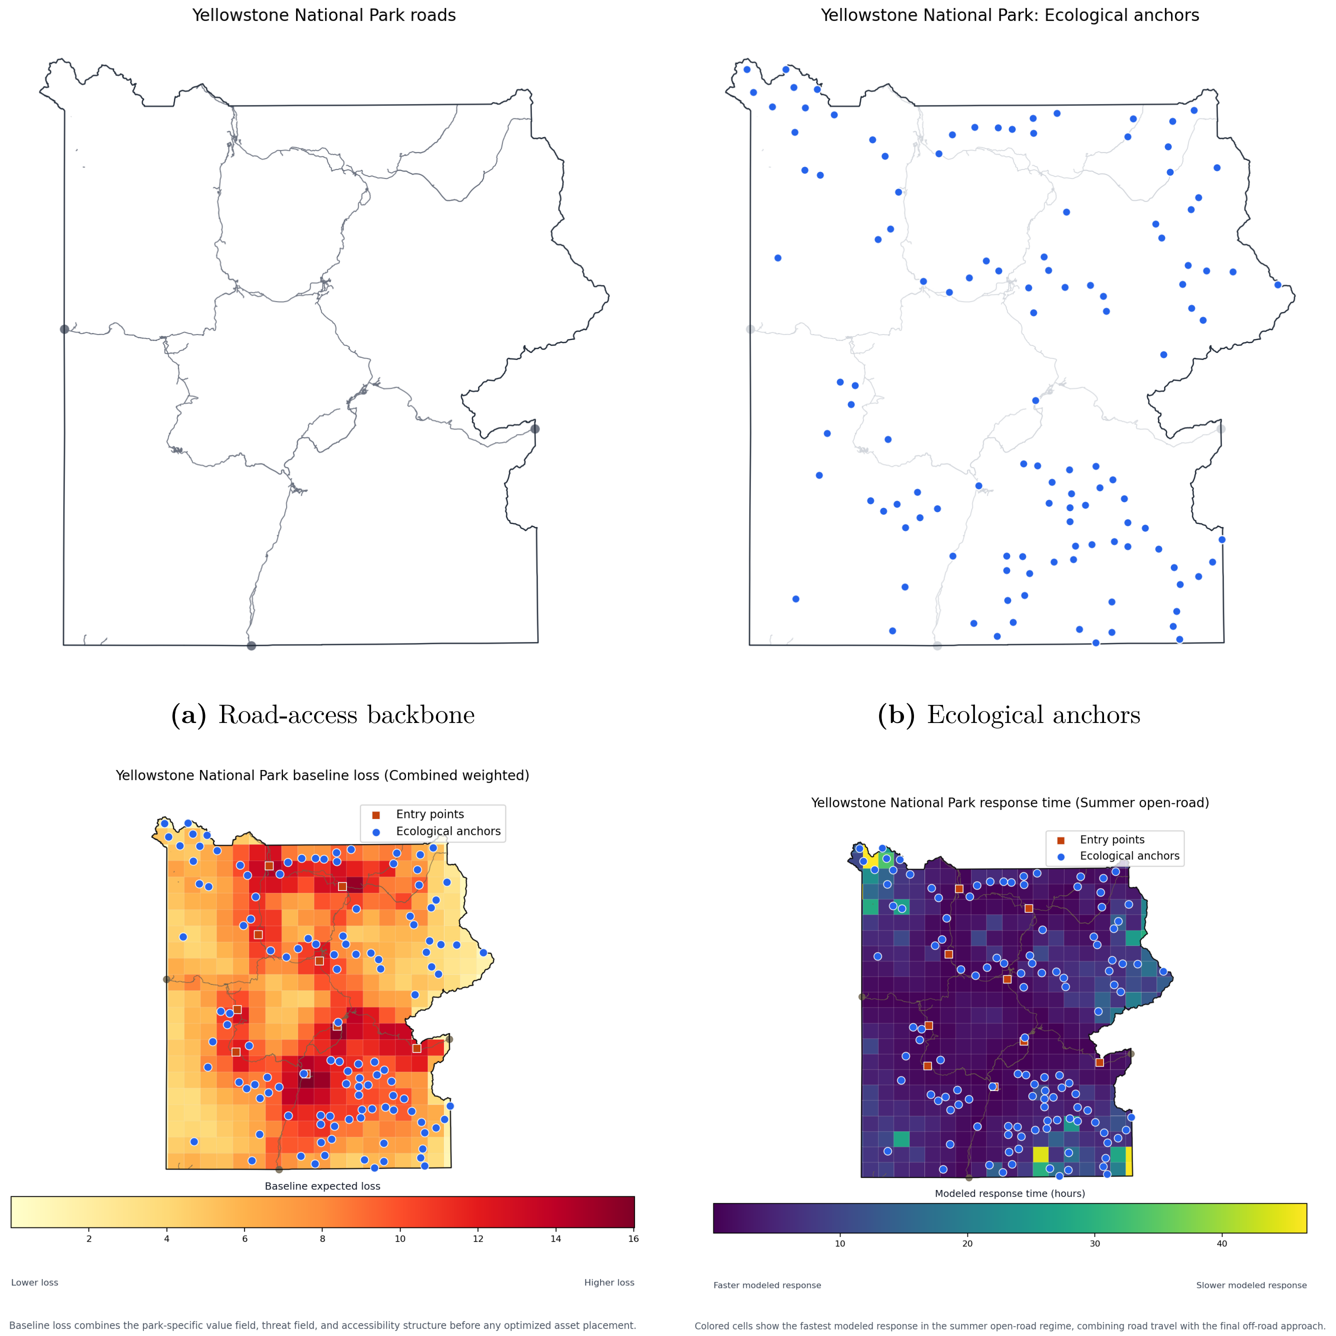

(a) Road-access backbone

(b) Ecological anchors

(c) Combined baseline loss

(d) Open-road response time

Figure 15: Yellowstone empirical rerun. The transferred pipeline rebuilds a road-driven and visitor-sensitive geometry rather than an Etosha-style waterhole geometry. The response-time panel is shown in the summer open-road regime and is computed against the best reachable road component before adding the DEM-adjusted final off-road approach, making it a cleaner realism check on graph-constrained reach.

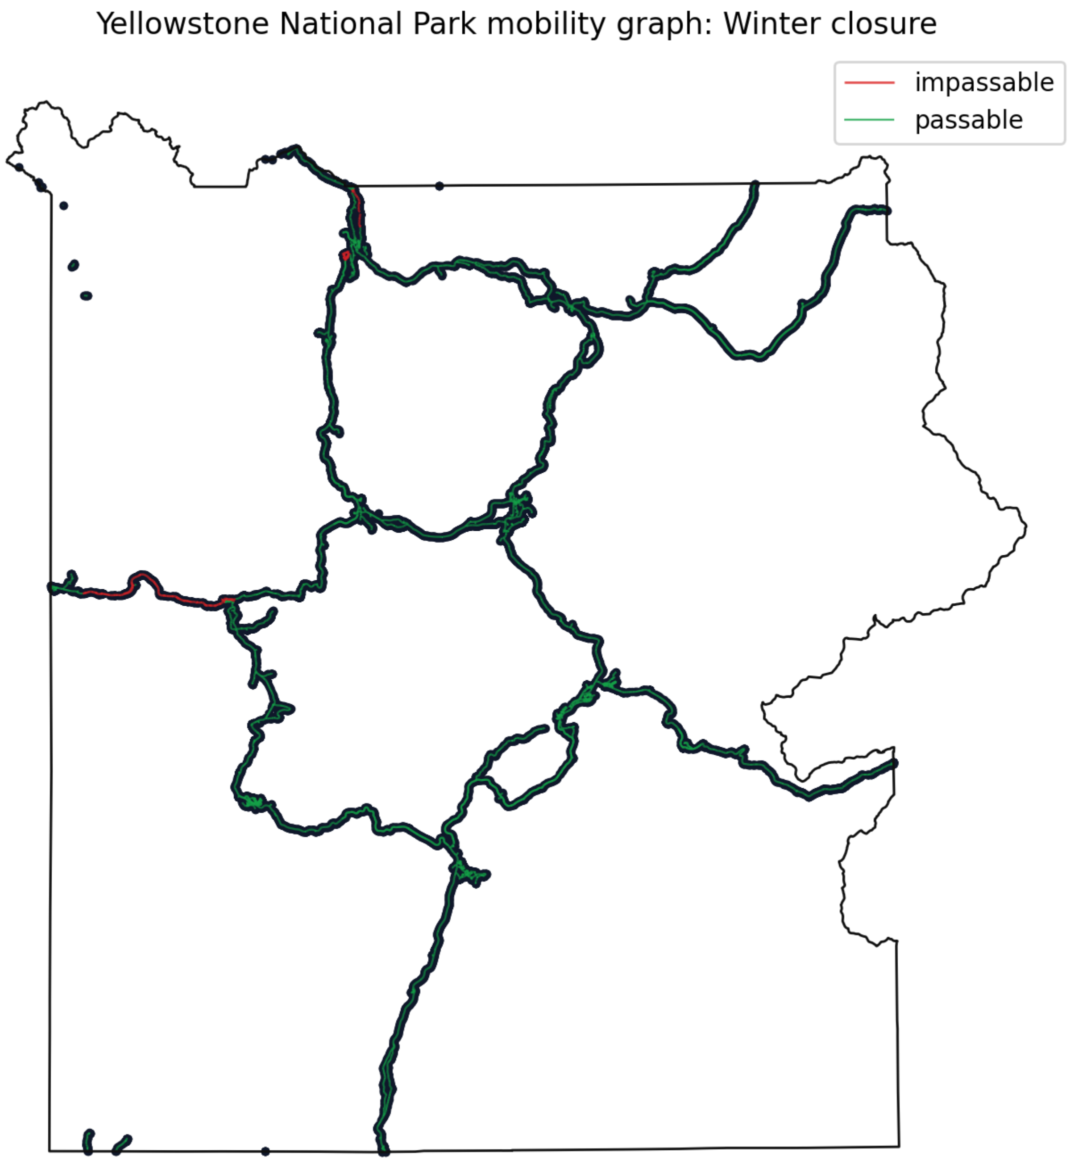

Figure 16: Yellowstone dynamic mobility graph. Passable edges remain green and winter-severed edges appear red. This is the key realism check for the Yellowstone rerun: the model is explicitly deleting seasonally constrained traversal edges before recomputing the response matrix.

Kaziranga rerun. Kaziranga is the complementary stress test because it keeps anti-poaching relevant but replaces static access with floodplain disruption. The rerun uses 4 operational access bases, of which 1 is an explicit OSM gate and 3 are supplementary south-boundary access points inferred from boundary-touching roads on the same OSM backbone. It also uses 11 crossing nodes, 8 refuge nodes, 239 hydrology points, 219 corridor segments, and an OpenTopography DEM layered onto the floodplain road-river structure. This immediately breaks the Etosha logic in a different direction: the important geometry is no longer scarce dry-season water, but flood-sensitive crossings, corridor chokepoints, and elevated refuge structure. The combined baseline loss map therefore becomes elongated and corridor-heavy. The non-flood response surface now shows a plausible multi-range access gradient across the park, while the flood-sensitive edge-deletion graph still makes the seasonal collapse in access explicit once inundation-sensitive edges are removed. Quantitatively, the non-flood median/p95 total response is about 0.88 and 1.95 hours, respectively, while the flood-sensitive regime shifts those to about 1.31 and 2.16 hours, respectively. Median road-only travel remains below a quarter-hour in both regimes, so the seasonal penalty is coming from flood-constrained final reach. Consult the figures below.

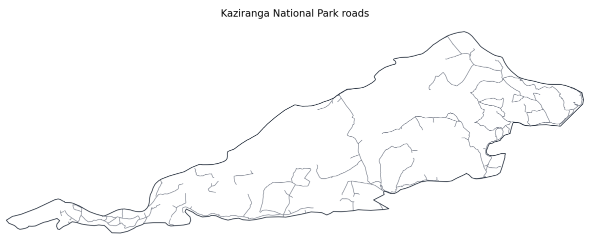

(a) Road-access backbone

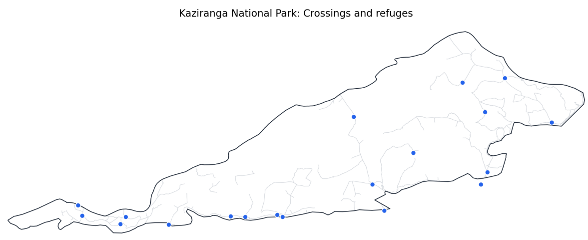

(b) Crossings and refuges

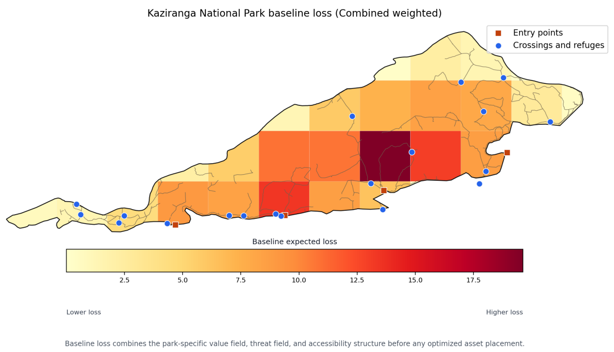

(c) Combined baseline loss

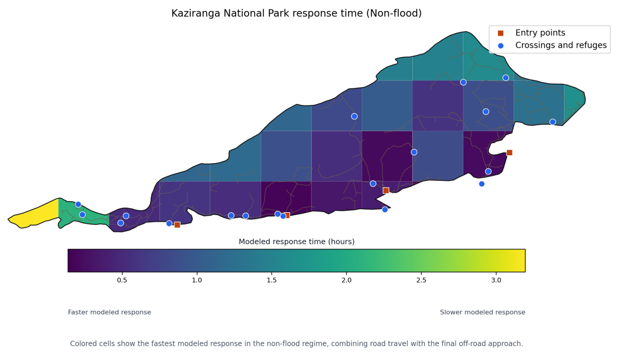

(d) Non-flood response time

Figure 17: Kaziranga empirical rerun. The model now centers on corridor, crossing, and refuge geometry rather than waterhole aggregation (Etosha). The response-time panel is shown in the non-flood regime using the corrected reachable-road logic and the DEM-adjusted final approach term, so it functions as the right companion to the flood-sensitive edge-deletion graph: baseline reach is broad but still structured, and seasonal inundation then squeezes the network sharply.

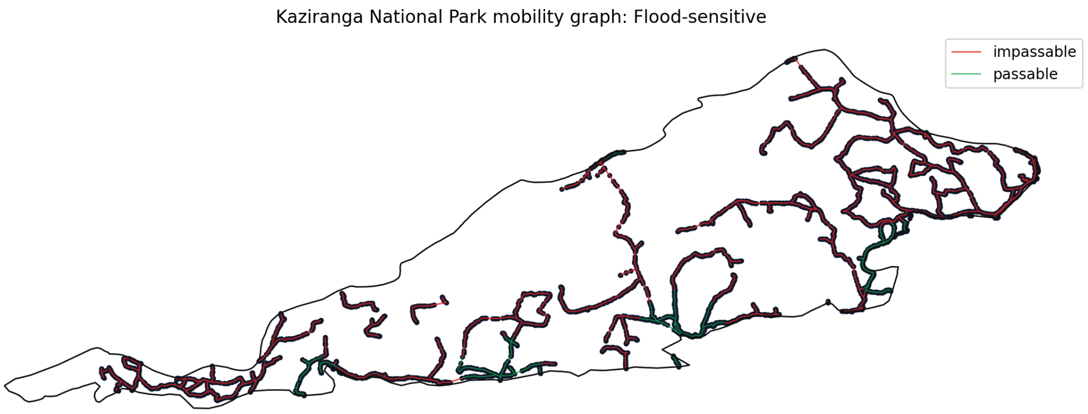

Figure 18: Kaziranga dynamic mobility graph. The flood-sensitive regime deletes quite a large share of low-lying traversal edges, leaving a visibly compressed accessible network. We believe that is exactly the kind of structural change that a transferable reserve-protection model must survive if it is to be more than an Etosha-specific road map.

Taken together, these reruns show the excellent transferability of the framework. The same loss-based objective, hybrid grid-graph geometry, terrain-aware response-time, and asset-allocation engine remained intact, despite the park-specific features changing sharply. In particular, Yellowstone replaced waterhole concentration by visitor-habitat and winter-closure structure. Kaziranga replaced a stable road graph by a corridor-and-flood regime. The fact that both reruns completed, produced coherent baseline loss and response surfaces, and did so in separate output trees without disturbing the Etosha outputs is exactly what we wanted in here, Appendix I.

§11. Appendix J. Core Code Snippets

The full codebase contains many thousands of lines, so this appendix does not attempt to reproduce everything. Instead, it collects the most implementation-critical snippets: terrain ingestion, terrain-adjusted mobility, and the final Pyomo allocation model. These are the sections that most directly support the paper's main claims about portability, realistic response, and optimization.

§11.1 DEM Ingestion from OpenTopography

The first transferability claim is that terrain can be regenerated automatically for a new park. Listing 1 is the DEM data from `src/py_optim/dem_tools.py`. It analyzes the park boundary, expands the bounding box slightly, and requests a GeoTIFF DEM from OpenTopography using the configured API key.

Listing 1: Commented excerpt of the OpenTopography DEM fetch pipeline

```
def fetch_dem_if_missing(cfg, boundary_gdf, *, force=False, logger=None):
    existing = find_input_file(cfg.raw_dir,("dem","etosha_dem"), RASTER_EXTENSIONS)
    target_path = cfg.raw_dir/"dem.tif"
    if existing is not None and not force:
        return existing

    # Only opentopography is implemented here.
    provider = str(cfg.get("terrain", "dem_provider",
        default="opentopography")).strip().lower()
    if provider != "opentopography":
        return existing

    api_key_env, api_keys = _candidate_opentopography_keys(cfg)
    if not api_keys:
        return existing

    # Convert park boundry to lat/lon & add small buffer so
    # terrain raster slightly bigger than polygon.
    minx, miny, maxx, maxy = boundary_gdf.to_crs("EPSG:4326").total_bounds
    buffer_deg = float(cfg.get("terrain", "dem_bbox_buffer_deg", default=0.01))
    south, north = max(-90.0, miny- buffer_deg), min(90.0, maxy + buffer_deg)
    west, east = max(-180.0, minx - buffer_deg), min(180.0, maxx + buffer_deg)

    params = {
        "demtype": str(cfg.get("terrain", "dem_product",default="COP30")),
        "south": south,
        "north": north,
        "west": west,
        "east": east,
        "outputFormat": "GTiff",
        "API_Key": api_keys[0],
    }

    response = requests.get(
        str(cfg.get("terrain", "opentopography", "base_url")),
        params=params,
        timeout=int(cfg.get("terrain" , "dem_timeout_s", default=600)),
    )
    response.raise_for_status()
```

```
    target_path.write_bytes(response.content)
    return target_path
```

§11.2 Terrain-Adjusted Travel Speeds

Fetching the DEM is only useful if it actually changes the movement model. Listing 2 shows the implemented terrain handling: reprojecting the DEM into the park CRS when needed, computing slope in degrees, sampling slope at analysis points, and converting slope into off-road and on-road speed penalties. This is the code that turns "terrain-aware" from a claim into a numeric part of the model.

Listing 2: Commented excerpt of terrain loading and slope-adjusted speed functions

```python
def load_dem_context(cfg, projected_crs, *, logger=None):
    dem_path = find_input_file(cfg.raw_dir, ("dem", "etosha_dem"), RASTER_EXTENSIONS)
    if dem_path is None:
        return None

    with rasterio.open(dem_path) as src:
        # Reproject on the fly so slope and distance are computed in the
        # same projected CRS used by the park grid and road graph.
        if src.crs is None or str(src.crs) != projected_crs:
            with WarpedVRT(src, crs=projected_crs, resampling=Resampling.bilinear) as vrt:
                data = vrt.read(1, masked=False).astype(float)
                nodata = vrt.nodata
                transform = vrt.transform
        else:
            data = src.read(1, masked=False).astype(float)
            nodata = src.nodata
            transform = src.transform

    if nodata is not None:
        data[data == nodata] = np.nan
    slope_deg = compute_slope_degrees(data, transform)
    return DemContext(path=dem_path, crs=CRS.from_user_input(projected_crs),
                      transform=transform, data=data, slope_deg=slope_deg, nodata=nodata)

def tobler_speed_kmh(slope_deg):
    # Tobler walking function for off-road terminal travel.
    slope_array = np.asarray(slope_deg, dtype=float)
    tan_theta = np.tan(np.radians(slope_array))
    return 6.0 * np.exp(-3.5 * np.abs(tan_theta + 0.05))

def road_speed_factor(slope_deg, cfg):
    # Gentle penalty for driving on steeper road segments.
    correction = float(cfg.get("terrain", "road_slope_correction", default=0.18))
    min_factor = float(cfg.get("terrain", "min_road_speed_factor", default=0.65))
    tan_theta = np.abs(np.tan(np.radians(np.asarray(slope_deg, dtype=float))))
    factor = 1.0 - correction * tan_theta
    return np.clip(factor, min_factor, 1.0)
```

§11.3 Dynamic Mobility Graph and Reachable-Component Response

The next claim that required code backing was realistic response time. Listing 3 shows the implemented response-graph logic from `src/py_optim/build_graph.py` in our repository. In the code the first block applies park-specific edge deletion for Yellowstone winter closures or Kaziranga flood masking. The second block computes the best reachable response through active graph components rather than naively snapping each cell to one nearest road stub.

Listing 3: Commented excerpt of seasonal edge deletion and reachable-component response logic

```
def _regime_constraints(cfg, *, prepared, graph_edges, bases, logger):
    constraints = {}

    if park_id(cfg) == "yellowstone":
        closures = load_context_layer(cfg, "winter_closures",
                                      projected_crs=prepared.metadata["projected_crs"],
                                      logger=logger)
        closure_geom = closures.geometry.union_all()
        constrained = closure_geom.buffer(float(cfg.get("park_model",
            "closure_buffer_m", default=200.0)))
        edge_mask =
            ~graph_edges.geometry.intersects(constrained).to_numpy(dtype=bool)
        base_mask = ~bases.geometry.intersects(constrained).to_numpy(dtype=bool)
        constraints["regime_2"] = (edge_mask, base_mask, 0.85)

    if park_id(cfg) == "kaziranga":
        flood_mask = load_context_layer(cfg, "flood_mask",
                                        projected_crs=prepared.metadata["projected_crs"],
                                        logger=logger)
        flood_geom = flood_mask.geometry.union_all()
        edge_mask =
            ~graph_edges.geometry.intersects(flood_geom).to_numpy(dtype=bool)
        base_mask = ~bases.geometry.intersects(flood_geom).to_numpy(dtype=bool)
        constraints["regime_2"] = (edge_mask, base_mask, 0.80)

    return constraints

# For each scenario, solve shortest travel from all active bases at once.
travel_to_node = nx.multi_source_dijkstra_path_length(graph, base_graph_ids,
    weight="travel_h")
component_ids, component_by_node = _component_index(graph, node_ids)
active_component_ids = sorted({
    component_by_node.get(base_graph_id, -1)
    for base_graph_id in base_graph_ids
    if component_by_node.get(base_graph_id, -1) >= 0
})
component_node_indices = {
    component_id: np.flatnonzero(component_ids == component_id)
    for component_id in active_component_ids
}

for cell_id, cell_x, cell_y in zip(centroid_ids, centroid_x, centroid_y,
    strict=False):
    best_response_h = float("inf")
    for component_node_index in component_node_indices.values():
```

```
        distances = np.hypot(node_x[component_node_index] - cell_x,
                          node_y[component_node_index] - cell_y)
        candidate_count = min(len(component_node_index), candidate_nodes_per_component)
        candidate_local_indices = np.argpartition(distances, candidate_count - 1)[:candidate_count]
        for local_idx in candidate_local_indices:
            node_idx = int(component_node_index[int(local_idx)])
            road_h = float(travel_to_node.get(str(node_ids[node_idx]), np.inf))
            offroad_h = (float(distances[int(local_idx)]) / 1000.0) / max(offroad_speed, 0.5)
            response_h = road_h + offroad_h
            if response_h < best_response_h:
                best_response_h = response_h
```

§11.4 Core Allocation MILP

Finally, Listing 4 shows the allocation model itself from `src/py_optim/optimize.py`. This is the executable core of the paper's loss-minimization logic: asset-selection variables, monitoring aggregation, piecewise-linear detection, budget constraints, and the objective of avoided-loss.

Listing 4: Commented excerpt of Pyomo allocation model

```
model = pyo.ConcreteModel(name=f"{park_short_name(cfg)}_{scenario_name}")
model.G = pyo.Set(initialize=cell_ids) # grid cells
model.R = pyo.Set(initialize=route_ids) # patrol-route candidates
model.D = pyo.Set(initialize=sortie_ids) # drone-sortie candidates
model.Sites = pyo.Set(initialize=site_ids) # sensor-site candidates

# Binary deployment decisions: choose or do not choose each asset.
model.y = pyo.Var(model.R, domain=pyo.Binary)
model.z = pyo.Var(model.D, domain=pyo.Binary)
model.x = pyo.Var(model.Sites, domain=pyo.Binary)

# M[g] is aggregate monitoring intensity in cell g.
# S[g] is the piecewise-linear proxy for detection success in cell g.
model.M = pyo.Var(model.G, domain=pyo.NonNegativeReals)
model.S = pyo.Var(model.G, bounds=(0.0, max(intercept for _, intercept in lines)))

def monitoring_rule(model, g):
    return model.M[g] == (
        sum(phi_map.get((r, g), 0.0) * model.y[r] for r in model.R)
        + sum(psi_map.get((d, g), 0.0) * model.z[d] for d in model.D)
        + sum(chi_map.get((s, g), 0.0) * model.x[s] for s in model.Sites)
    )

model.monitoring = pyo.Constraint(model.G, rule=monitoring_rule)

# Piecewise-linear upper envelope: detection rises with monitoring,
# but with diminishing returns according to calibrated breakpoints.
model.piecewise = pyo.ConstraintList()
for g in model.G:
    for slope, intercept in lines:
        model.piecewise.add(model.S[g] <= intercept + slope * model.M[g])

```

```
# Resource ceilings: staffing, drone count, sensor count, and hours.
model.patrol_budget = pyo.Constraint(
    expr=sum(route_staff[r] * model.y[r] for r in model.R) <= patrol_budget
)
model.drone_budget = pyo.Constraint(
    expr=sum(sortie_units[d] * model.z[d] for d in model.D) <= drone_budget
)
model.sensor_budget = pyo.Constraint(
    expr=sum(model.x[s] for s in model.Sites) <= sensor_budget
)
model.patrol_hours = pyo.Constraint(
    expr=sum(route_hours[r] * model.y[r] for r in model.R) <= patrol_hours_budget
)
model.drone_hours = pyo.Constraint(
    expr=sum(sortie_hours[d] * model.z[d] for d in model.D) <= drone_hours_budget
)

# Optional floor constraints used in robustness and fairness analysis.
if hotspot_cells:
    model.hotspot = pyo.Constraint(
        expr=sum(model.M[cell_id] for cell_id in hotspot_cells)
        >= hotspot_min_monitoring * len(hotspot_cells)
    )
if remote_cells:
    model.fairness = pyo.Constraint(
        expr=sum(model.M[cell_id] for cell_id in remote_cells)
        >= fairness_min_monitoring * len(remote_cells)
    )

# Objective: maximize expected avoided loss.
if scenario_name == "combined_weighted":
    model.objective = pyo.Objective(
        expr=sum(
            sum(
                combined_component_loss[name, g]
                * combined_response_prob[name, g]
                * model.S[g]
                for g in model.G
            )
            for name in ("regime_1", "regime_2", "intrusion", "fire")
        ),
        sense=pyo.maximize,
    )
else:
    model.objective = pyo.Objective(
        expr=sum(base_loss_map[g] * response_map[g] * model.S[g] for g in model.G),
        sense=pyo.maximize,
    )
```

Taken together, these code snippets above represent the most important bridges between the working software and our mathematical model. Listing 1 is the code showing terrain is reproducibly fetched instead of added manually. Listing 2 shows how terrain changes movement including its code representing vehicle and hiking speed. Listing 3 is the seasonal disruptions effect to response graphs.

Finally, Listing 4 shows how exactly the final deployment is selected under budget.

§12. Appendix K. Planning Calibration Guide

This appendix K is the extensive list/record of the planning calibration choices that were used in our present Etosha model. Moreover, it includes why they are defensible for this park. Appendix E earlier is a compact parameter-grounding guide; Appendix K is the broader planning rationale and portability guide. These are not universal constants, and are best read as the current Etosha calibration: values chosen to reflect the park's waterhole-centered ecology, road-access geometry, staffing context, and available public data layers.

In another protected area/national park, the same mathematical structure can transfer, but the numbers in this appendix should be re-estimated, locally.

Table 24: Planning calibration guide

Calibration family	Present Etosha values	Why this is fair for Etosha and what should change elsewhere
Grid discretization	5.000 km clipped cells	A 5.000 km grid resolves waterhole belts and camp–gate corridors while still allowing repeated MILP solves and stress tests. Other parks should adjust grid size to area, access density, and hotspot scale.
Road and off-road movement	road speeds 45/30 km/h; off-road 5 km/h; slope correction 0.18; min factor 0.65	These values we believe are reasonable and reflect slower reserve travel due to observation stops, wildlife caution, and road quality, while off-road motion is slower and terrain-dependent (Haklay & Weber, 2008; OpenTopography, 2021; Tobler, 1993). Other parks of course should adjust for pavement, snow, rivers, or floods.
Visibility geometry	observer 6 m; target 1.7 m; step 250 m	The visibility model targets strategic sensing rather than fine optics. Heights reflect elevated platforms and human-scale targets, and the step balances speed and realism (OpenTopography, 2021). Recalibrate for vegetation or platform differences.
Ecological distance-decay lengths	dry-water 12 km; wet-water 24 km; dry-elephant 18 km; wet-elephant 28 km; dry-herbivore 22 km; wet-herbivore 35 km; fence 65 km; fire-water 30 km	The key pattern is shorter dry-season decay, forcing concentration near permanent water. This fits Etosha's dryland ecology (*Map of Etosha*, n.d.). Other parks should recalibrate based on habitat or flood/snow structure.
Access and fire distance decays	fence 12 km; gate 18 km; road 10 km; intrusion 15/10/12 km; fire 28/25 km	These encode how access risk fades with distance. They reflect Etosha's gate, road, and boundary structure plus fire context (NASA, n.d.-a). Other parks should adjust for permeability and barriers.

Calibration family	Present Etosha values	Why this is fair for Etosha and what should change elsewhere
Consequence weights	dry: rhino 1.0, elephant 0.6, herbivore 0.3; wet: rhino 1.0, elephant 0.6, herbivore 0.25; fire 1.0	These are priority weights, not biological exchange rates. They reflect IMMC's emphasis on high-value species. Other parks should change both species and ordering - it is undoubted that this changes dramatically.
Hazard weights	poaching 0.45/0.20/0.20/0.15; intrusion 0.50/0.30/0.20; fire 0.55/0.25/0.20	These reflect relative importance of wildlife attraction, access, and exposure. Exact values should be recalibrated using local incident knowledge.
Patrol-library geometry	speed 18 km/h; stop 0.35 h; max 2 routes/base; radii 40/28 km; 4 perimeter routes; decay 10 km; scale 1.5; min 0.02	Produces a realistic route library without over-expansion. Gate patrols are more local, and decay keeps influence corridor-shaped. Other parks should rebuild around their own bases and patrol doctrine.
Drone-library geometry	max 12 sorties; spacing 14 km; duration 2 h; radius 16 km; scale 1.8; min 0.02	Drones have broader but time-limited reach. Values reflect strategic assumptions. Adjust for canopy, aviation limits, or launch constraints.
Sensor siting and detection	12 waterhole sensors; spacing 12 km; 4 interior sites; spacing 18 km; radius 12 km; base 0.95; decay 8 km; min 0.08	Sensors are persistent but local. Line-of-sight and distance still matter, and waterholes/camps anchor placement. Other parks should adjust for vegetation and maintenance access.
Baseline optimization budgets	patrol 10; drone 4; sensor 12; patrol-hours 80; drone-hours 24	These define the baseline deployment package. These are – we believe, optimized – planning choices that are aligned with our resource-constrained setting. Other parks should reset to real capacity.
Response-success calibration	$\tau_0 = 2.000\,\mathrm{h}$; $\kappa = 2.2$	The midpoint is the practical threshold in which response becomes ineffective. Chosen for interpretability, not as measured data. Other parks, for implementation, should sensitivity-test this breakpoint.
Detection-envelope calibration	breakpoints $(0, 1, 2, 4, 6)$; values $(0, 0.12, 0.22, 0.34, 0.45)$	Encodes diminishing returns while keeping the MILP concave and probability-safe. The cap ensures detection remains partial. Recalibrate if better empirical data exists.
Hotspot and fairness settings	hotspot on; quantile 0.90; min 0.20; fairness off	Makes the baseline explicitly hotspot-first while leaving fairness as a policy choice. Turn fairness on if minimum service is required.

Calibration family	Present Etosha values	Why this is fair for Etosha and what should change elsewhere
Combined-scenario weights	dry 0.40; wet 0.35; intrusion 0.15; fire 0.10	These are managerial weights. They center the scenario on poaching while retaining other risks. Adjust for different threat profiles, but we believe this is reasonable for Etosha upon our research.
Staffing conversion	patrols $\{4, 6, 8, 10, 12, 14, 16\}$; 8 h/team; team sizes 6/2; sensors 1 per 4; base 12; targets 0.12–0.16	Translates activity into staff equivalents for planning. Designed for interpretability. Other parks should adapt to real staffing structures.
Stress-test multipliers	poaching 1.25; fire 1.35; staff 0.70; drones 2; sensors 8; response 1.30	Designed to expose structural weaknesses without being unrealistic. Other parks should tailor stress scenarios to local failure modes.

Transfer checklist

The key point of this appendix is not that Etosha's current numbers used in our mathematical model are necessarily universally correct. It is that they are transparent, park-specific, and mathematically interpretable. The structure of the model transfers; the calibration does not. We would like to point out that any serious deployment in another park should preserve the loss-based architecture while re-estimating the values in this appendix from local ecology, access, staffing doctrine, and incident history.