

Bridging the gaps: Agent-based modelling for elephant poaching mitigation

ODD Protocol

EMILY NEIL, JENS KOED MADSEN, ERNESTO CARRELLA, NICOLAS PAYETTE, RICHARD BAILEY

Purpose

As an illustration of the utility of agent-based models (ABMs) for poaching mitigation, we developed an exploratory ABM that predicts how interactions between elephants, poachers, and law enforcement affect poaching levels within a virtual protected area. The model is theoretical at this stage but is parameterized with realistic ecological and behavioral data on African elephants, as well as representative information on poaching and ranger strategies. The aim of this model is not to provide a realistic depiction of poaching, but instead to demonstrate how ABMs can bridge the gaps present in the other two main modelling techniques applied to this topic to-date (equation-based and game theoretical models) and to provide a framework for future research. The model provides a starting point for further development and application to real-world situations, perhaps incorporating real GPS data of elephant movements and poaching incidents, and GIS satellite imagery. The aim is for this model to be further developed into a useful management support tool, one that can be used as a virtual laboratory to experiment with different scenarios without putting time, funds, resources, personnel, or elephants at risk.

State Variables and Scales

The model landscape is a simplified representation of a protected area that is split into four ‘zones.’ This virtual park is populated by both elephants and poachers. The size of the landscape and the timing of events are arbitrary and do not coincide with a real-world situation. The environment is split into four zones, each of which receives a different amount of law enforcement effort. The entities in the model are described in Table 1.

Entity	Parameters	Range and Unit
Elephants	Initial number of elephants	150
	Number of herds	10
	Sex	M/F
	Age	1-60 years
	Status	Matriarch or follower
	Herd Number	Unique number
Poachers	Number of poachers	15
	Exploration probability	10%
	Profit memory	A set of values associated with hunting locations, based on where the poacher has previously seen elephants, and on where other poachers have been caught by law enforcement

	Hunting effectiveness	50%
Law Enforcement responses	Probability of catching poacher	0-50% per zone depending on management technique

Table 1 Attributes for the agents in the model

The temporal scale in the model is represented in days, with each time step (tick) equaling one day. The spatial extent of the model is 40 x 40 patches, and this is then split into four zones (1-4). Patch size is not specified, and thus does not represent specific spatial dimensions. However, it is assumed to be consistent with expected protected area size requirements to support multiple herds of elephants.

Process Overview and Scheduling

Each day, the following processes are called in the given order. State variables are updated immediately. The sub-models implementing these procedures are described in-detail in Section 3.10.

1. *Elephant Migration & Dispersal*: Elephant matriarchs migrate to different zones three times a year, and the rest of the herd follows her. Elephants with the status ‘follower’ follow their matriarchs, while also maintaining close proximity to other elephants in their herd. Male elephants above the age of fourteen disperse from their matriarchal herds and begin to migrate independently.
2. *Reproduce*: Female elephants reproduce every five years if they are thirteen years of age or older.
3. *Poacher Movement*: Poachers start the simulation in the village, with a random number of days (between 1-10) that they will stay in the village. When the countdown reaches 0, they will start a poaching trip. Depending on the scenario being tested, they will either move randomly across the landscape to hunt elephants or they will move according to an epsilon-greedy bandit algorithm. They will keep track of which zones they have seen elephants in, and which zones other poachers have been caught in, and update their beliefs about the ‘best’ zones accordingly.
4. *Catch Elephants*: If a poacher moves to a patch during a poaching trip and an elephant is there, the elephant will ‘see’ the elephant and remember which zone it was in. They then have a probability of effectively catching and killing the elephant (poaching effectiveness; set at 50% in this model).
5. *Law Enforcement Schedule*: Depending on the scenario being tested, the probability of catching a poacher per zone varies between 0-50%. In scenarios 1.1, 1.2, and 2.1, law enforcement effort is distributed unevenly across the entire protected area, at a rate of once per week. This is to simulate law enforcement scheduling, as they may not be patrolling every day of the week. In scenario 2.2, rangers patrol by adaptively following elephant matriarchal herds, and the probability of catching a poacher is therefore highest where the matriarchal herds are located.

6. *Death*: Elephants can die from old age or poachers. Poachers exit the system (“die”) if they are caught by law enforcement.

Design Concepts

Adaptation

Poacher agents move from the village into one of the four zones to hunt for elephants. Depending on the scenario being tested, poachers will either move randomly across the four zones, or they will dynamically adapt to elephant whereabouts and various law enforcement interventions using an epsilon-greedy bandit algorithm (Sutton & Barto, 1998). When moving adaptively, poachers have an exploration probability of 100% for the first five trips to gain an understanding of the different zones (e.g. where elephants are located and where they are more likely to be caught by law enforcement). After this exploration period, they explore a zone at random according to probability epsilon ($\epsilon = (0, 1)$), or otherwise return to the zone that had the ‘best’ outcome (e.g. the zone in which they saw the most elephants, and in which the fewest poachers were caught by law enforcement) (Kuleshov & Precup, 2014; Sutton & Barto, 1998). In other words, poachers face an exploitation-exploration trade-off: they must choose between hunting in the best zone or continuing to learn about the system (Bubeck, 2012; Kuleshov & Precup, 2014). Individual poachers thus learn which zone is most profitable to poach in as a consequence of theirs and other poachers’ experiences: they remember how many elephants they have personally seen in each zone, and they remember in which zones poachers have been most frequently caught by law enforcement. The hunting effectiveness (how often poachers successfully kill an elephant once they target it) can be varied in the model, as different types of weapons and technologies will change the success rates of poachers. Poachers caught by law enforcement permanently leave the system.

Interaction

Poachers interact with law enforcement and elephants to learn more about the ‘best’ zone in the protected area. Depending on the scenario being tested, law enforcement has abstracted interactions with elephant herds, as they will follow matriarchal herds on their patrol. Law enforcement also interacts with poachers, catching them and removing them from the system.

Sensing

Poachers are assumed to immediately have access to all information regarding the location of other poachers who have been caught by law enforcement. They use this information to update their beliefs about which zone is the best to hunt in.

Implementation Details

Initialization

The model was implemented in NetLogo version 6.0 (Wilensky, 1999). The landscape, elephants, and poachers are initialized when the model starts. Elephants are initialized by

creating 150 individuals and 10 herds. Their location is set randomly within the protected area. Elephants are randomly assigned sex, an age, a herd number, and a status within the herd. Poachers are initialized by creating 15 individuals, with a hunting effectiveness of 50%, and an exploration probability of 10%.

Sub-models

Elephant Migration and Dispersal

Two types of elephant movement were simulated: 1. herd movement, and 2. adult males dispersing. For herd movement, all female elephants and males below the age of 14 with the same herd number follow the migratory movements of their matriarch. Matriarchs follow a migration pattern, moving to a new area of the protected area three times a year. Migratory routes do not change from year to year in the model, and different elephant herds can overlap and feed on the same optimal patches at the same time.

The model reflects known elephant behavior in that herds usually consist of adult females – led by a matriarch - and their immature offspring. Female and young male elephants follow a matriarch and move as a herd, following a seasonal migratory pattern as they would in reality (Thouless, 1995). Adult males (>14 years old) disperse from the matriarchal herds and move independently (Moss, 2000).

Reproduction

Female elephants in the model reproduce once every five years if they are above the age of thirteen. This follows known elephant ecology: female elephants start to breed at around eleven, producing their first calf at around thirteen years of age, and only give birth every four to five years on average (Moss, 1988).

Death

Elephants die of old age (>60 years) or being caught by poachers. The oldest female in a group is the matriarch, and if the matriarch dies, the next eldest female takes over the role. Poachers “die” - leave the system - when caught by law enforcement.

Poacher Movement

Poachers leave the village to hunt. Depending on the scenario being tested, they either move randomly from zone to zone or they move adaptively according to elephant whereabouts and law enforcement strategies. When moving adaptively, poachers have an exploration probability of 100% for the first five ticks of the simulation so they may gain an understanding of the different zones (where elephants are located and where they are more likely to be caught by rangers). After this exploration period, they weigh the exploitation-exploration tradeoff to choose which zone to poach in.

Catch Elephants

If a poacher finds an elephant when moving zones, the poacher sees the elephant and has a 50% probability (hunting effectiveness) of killing it. The hunting effectiveness can be varied in the model, as different types of weapons and technologies will change the success rates of poachers.

Law Enforcement Schedule

We first compared a scenario in which poacher agents are non-adaptive (scenario A) to one in which poachers dynamically adapt to law enforcement strategies and elephant whereabouts (scenario B). In both scenarios, law enforcement effort is distributed unevenly across the protected area. Each of the four zones has a different probability of catching a poacher (ranging from 0-25% for a total of 50% over the entire protected area); this is often the case in reality, as some regions of a protected area are better covered by law enforcement than others. There is very little empirical data available on the probability of catching poachers and 50% may not reflect reality. Scenario 1 shows how ABMs can expand upon equation-based approaches by allowing for dynamic and adaptive poachers.

We then considered the effect of using elephant behaviour and ecology to inform law enforcement strategies. In scenario B, rangers patrolled in an uneven distribution across the protected area. In scenario C, rangers follow elephant matriarchal herds while patrolling. Poachers are adaptive in both scenarios. This experiment shows how ABMs can build upon game theoretical approaches by incorporating the behavior and ecology of elephants, and by opening up the possibility of exploring new management techniques outside of planning optimal patrol routes.

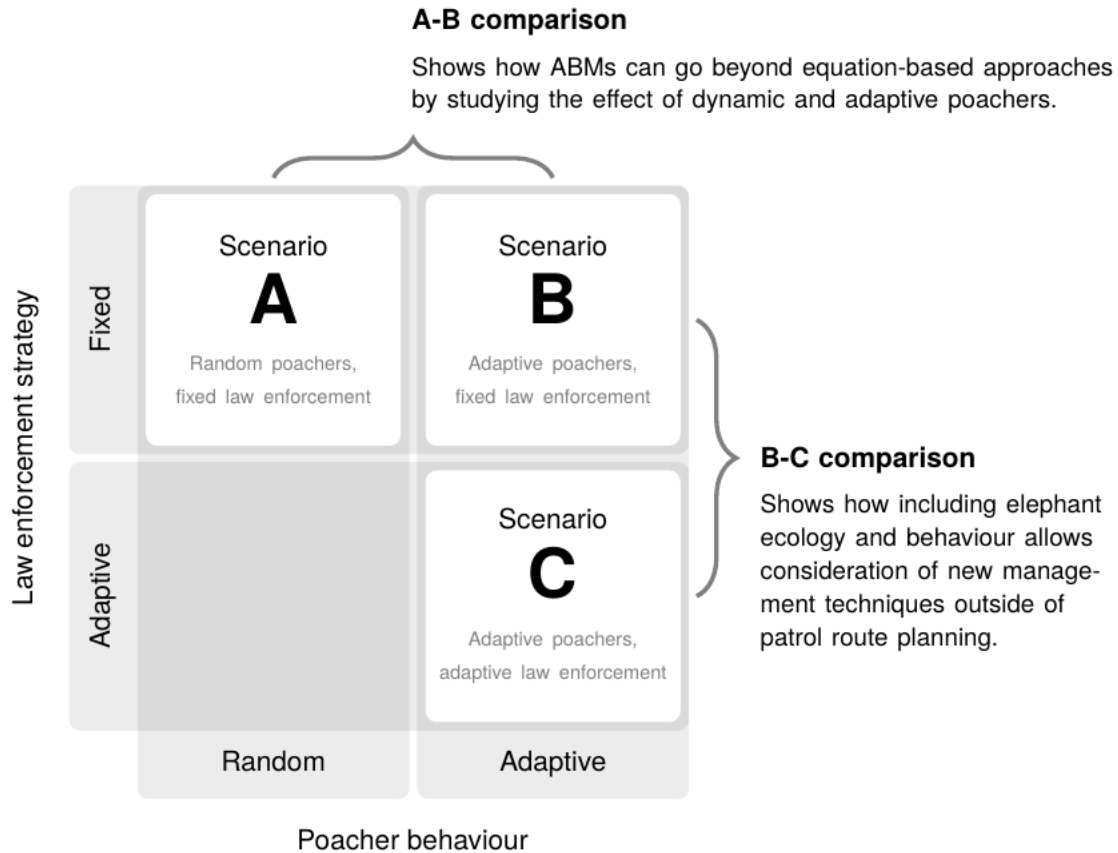


Figure 2: Descriptions of the three scenarios explored in the model

Data Analysis

The model was run in BehaviorSpace, a Netlogo tool that can run many simulations of a model and vary the settings of interest, and then records the results of each iteration (Wilensky, 1999). We ran each simulation 394 times, as determined by a power calculation for t-test of means (Lipsey, 1990; Seri & Secchi, 2017), for one year (365 ‘ticks’). We simulated the scenarios described in Figure 2 and counted elephants and poachers on each day (“tick”) until either poachers or elephants reached 0. The data was analyzed using RStudio (version 1.0.136).

References

- Bubeck, S. (2012). Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Foundations and Trends® in Machine Learning*, 5(1), 1–122.
<https://doi.org/10.1561/22000000024>
- Kuleshov, V., & Precup, D. (2014). Algorithms for multi-armed bandit problems. *ArXiv:1402.6028 [Cs]*. Retrieved from <http://arxiv.org/abs/1402.6028>
- Lipsey, M. W. (1990). *Design Sensitivity: Statistical Power for Experimental Research*. SAGE.
- Moss, C. (2000). *Elephant Memories: Thirteen Years in the Life of an Elephant*. University of Chicago Press.

- Seri, R., & Secchi, D. (2017). How Many Times Should One Run a Computational Simulation? In B. Edmonds & R. Meyer (Eds.), *Simulating Social Complexity: A Handbook* (pp. 229–251). https://doi.org/10.1007/978-3-319-66948-9_11
- Sutton, R. S., & Barto, A. g. (1998). *Reinforcement Learning. An Introduction*. The MIT Press.
- Thouless, C. R. (1995). Long distance movements of elephants in northern Kenya. *African Journal of Ecology*, 33(4), 321–334. <https://doi.org/10.1111/j.1365-2028.1995.tb01042.x>
- Wilensky, U. (1999). *NetLogo: Center for Connected Learning and Computer-Based Modeling*. Northwestern University.