An Expert-Driven Causal Model of the Rhino Poaching Problem

Hildegarde Koen*[a,b], JP de Villiers[a,b], Henk Roodt[c], Alta de Waal[d]

[a]*Defence Peace Safety and Security (DPSS), Council for Scientific and Industrial Research (CSIR), PO Box 395, Pretoria 0001, South Africa,*
hkoen@csir.co.za, jdvilliers1@csir.co.za
[b]*Department of Electrical, Electronic and Computer Engineering, University of Pretoria, Private Bag X20, Hatfield, 0028, Pretoria, South Africa*
[c]*Waikato Institute of Technology, New Zealand,*
henk.roodt@wintec.ac.nz
[d]*Department of Statistics, University of Pretoria, Private Bag X20, Hatfield, 0028, Pretoria, South Africa,*
alta.dewaal@up.ac.za

Abstract

A significant challenge in ecological modelling is the lack of complete sets of high-quality data. This is especially true in the rhino poaching problem where data is incomplete. Although there are many poaching attacks, they can be spread over a vast surface area such as in the case of the Kruger National Park in South Africa, which is roughly the size of Israel. Bayesian networks are useful reasoning tools and can utilise expert knowledge when data is insufficient or sparse. Bayesian networks allow the modeller to incorporate data, expert knowledge, or any combination of the two. This flexibility of Bayesian networks makes them ideal for modelling complex ecological problems. In this paper an expert-driven model of the rhino poaching problem is presented. The development as well as the evaluation of the model is performed from an expert perspective. Independent expert evaluation is performed in the form of queries that test different scenarios. Structuring the rhino poaching problem as a causal network yields a framework that can be used to reason about the problem, as well as inform the modeller of the type of data that has to be gathered.

Keywords: rhino poaching, predictive modeling, insufficient data, expert knowledge

1. Introduction

The rhino poaching problem has reached alarming heights and has remained a challenging problem for the past few years. However, what was once the latest buzzword has now become just

---

*Corresponding author

another story in the newspaper. The reality is that rhino poaching is still on the rise, with no chance of decreasing any time soon.

At the end of 2008 the poaching count stood at 83 rhinos and rapidly increased until a total of 1,175 rhinos were poached in 2015. The biggest problem is in the Kruger National Park (KNP), which is home to the largest concentration of rhinos in South Africa (Emslie & Brooks, 1999). The KNP is situated on the border between South Africa, Zimbabwe and Mozambique, which makes it more difficult to deter and apprehend poachers.

Different approaches have been followed to mitigate the rhino poaching problem, but without much success. Most approaches are reactive in the sense that they are only used (or feasible) after a rhino has been poached. If a shot is heard the location of the shot can be triangulated, but by then it is too late. If a poacher is caught with a rhino horn he is incarcerated, but it is too late for the rhino. The approach presented in this paper is proactive instead of reactive: rangers wish to know that a poaching event is likely to occur well before it does. A causal model is developed to predict the area where a next poaching event might occur. Utilising a reduced patrol surface area as informed by the model results in rangers and other resources being allocated more effectively. Extending the previous work by the authors in Koen et al. (2014), a current perspective predictive model of the rhino poaching problem is extended by a team of domain experts. This effort further serves to develop the only expert-driven causal model of which the authors are aware of that applies to wildlife crime. This is also the only known predictive model for the rhino poaching problem.

At the start of the project, data was scarce and an expert approach was a pragmatic method to define and understand the problem, as well as to make predictions. Fortunately for the modeller, but unfortunately for the rhinos, the number of recorded poaching attacks have increased to such an extent that there exists data for a couple of thousand poaching attacks over several years. However, the expert-driven model remains useful in understanding the rhino poaching problem, as well as reasoning about it. Another complication associated with the data is that poaching attacks are scattered over an area roughly the size of Israel. The data sparsity problem remains if the park is divided into one square kilometre cells for the purposes of prediction, as is the case in this study. However, as more data becomes available, a model based on data will increasingly be preferred over an expert model to inform predictions. This paper describes the development of the expert model, as well as exploring the very realistic possibility of not having data in a wildlife crime problem in countries or parks with limited resources to record data.

1.1. Literature Review

Initially very little was known about the rhino poaching problem and the drivers influencing it. Numerous papers have been published concerning alternative management strategies to cope with the continuously rising cost of rhino poaching (Kahler & Gore, 2012), (Duffy & St John, 2013), (Biggs et al., 2013), and (Ferreira et al., 2014), but almost none on modelling the rhino poaching problem. The paper by Ferreira et al. (2015) uses poaching records as well as growth rates for both black and white rhinos from 2008 and 2010 to predict the population sizes and the growth rates of rhinos in the KNP in 2013. A similar paper is presented by Haas & Ferreira (2015) where the same type of information is used to predict when the rhino populations will start to decrease rapidly. This is calculated by using a special form of individual-based model (IBM) called the "agent/individual-based economic-ecological model" (Haas & Ferreira, 2015). The study by Burn et al. (2011) analyses the trends and patterns concerning illegal elephant killings, but they use Bayesian hierarchical modelling in contrast with our causal Bayesian network.

Recently, two papers have been published that model and analyse facets of rhino poaching by following different approaches. The paper by Park et al. (2015) presents a special case of spatio-temporal optimisation problems (an anti-poaching engine) where both poacher and rhino behaviour models are supplied to the engine as input. In the paper by Critchlow et al. (2015), the spatial patterns of various illegal types of wildlife activities are investigated.

Critchlow et al. (2015) focus on identifying drivers of different types of illegal activities, whereas this study only focuses on the poaching of white rhinos. Park et al. (2015) also attempt to safeguard rhinos, but they do not differentiate between black and white rhinos. The behaviour, vegetation preference, and population numbers of white rhinos are very different to that of black rhinos, thus it would not make sense to simply group the two subspecies together under the common heading of "rhinos".

Another similarity between this study and Park et al. (2015) is that both investigate ways to mitigate rhino poaching attacks through prediction for the safeguarding of the rhino population. Both this study and that of Critchlow et al. (2015) understand that the patterns of illegal activities are important in solving the poaching problem, and that the drivers of poaching are poorly understood and documented.

The maps of the game reserves or parks in all three studies are divided into grids with square cells of equal size. For this study a grid size of 5×5 kilometre was used initially, but later on

it was changed to 1×1 kilometre for evaluation purposes, whereas in Park et al. (2015) a grid size of 400×400 metres is used. The focus of this study is on the KNP situated in the north of South Africa, spanning an area of 19,485 km^2 (Ferreira et al., 2015), whereas the park used in the study of Park et al. (2015) is Olifants West in South Africa, spanning an area of about 14×10 kilometres. Critchlow et al. (2015) focus on the Queen Elizabeth Conservation Area in Uganda which at 1,978 km^2 is roughly 10 times smaller than the KNP. They divide the map of the Queen Elizabeth Conservation Area into 500×500 metres. Both this study and the study by Critchlow et al. (2015) identify areas of greatest risk in terms of animal poaching, although this study focuses on rhino poaching. The grid sizes of all three studies are in line with the size of the corresponding game reserves or parks.

Further similarities between Park et al. (2015) and this study are the assumptions that poaching attacks usually occur during twilight and at night. The belief that rhinos avoid busy areas is also shared by both studies. Both studies calculate the probability of a cell being under attack by a poacher and utilise models which have both spatial and temporal components. Park et al. (2015) divide their temporal dimension into hourly intervals and use a different spatio-temporal graph for each agent, whereas this study uses a different instantiation of the causal model for each cell.

The biggest difference between the three studies is in terms of the approaches used. Park et al. (2015) use a special case of spatio-temporal optimisation problems combined with behavioural models, as well as multi-variate regression, whereas this study uses a Bayesian Network (BN). Critchlow et al. (2015) explicitly model ranger patrol effort, whereas this study predicts the probability of poaching events.

Park et al. (2015) also rely on certain assumptions that decrease the study's likelihood to yield usable results in a different environment. Firstly, the authors know where the rhinos in Olifants West are, because 20 of the rhinos in question are collared and have been tracked by Global Positioning System (GPS) for two years. In our study, the location of the rhinos is unknown as the area is immense and the rhinos are not collared, thereby adding an extra layer of difficulty to the problem.

Another difference between the studies is that Park et al. (2015) use drones in combination with rangers for patrolling. If the study is to be adapted and used in other environments, it is thus assumed that drones are permitted in those parks. The authors state that in the Olifants West study they had one drone per ranger. One drone can cover a relatively small area, but even for a park as small as Olifants West six drones were needed. If their study is to be adapted to the KNP,

4

for example, there will be significant challenges such as obtaining approval for flying the drones, buying enough drones for the entire park (drones are very expensive), maintaining all the drones, *etcetera*.

Although it seems that BNs are not used in other wildlife crime applications, BNs have been used for urban crime applications. The paper by Gholami et al. (2015) presents a dynamic BN (DBN) that can be used to assign optimal patrol routes to patrol units, be it community policing or the police force. Urban crime is, according to the authors, opportunistic, and the criminals learn the movements of the patrol unit and use it to their advantage. The authors state that "...the criminals can adapt their strategy by seeking crime opportunity in less effectively patrolled locations." It is thus important that the criminals do not know the whereabouts of the patrol unit, just as is the case with the rhino poaching problem. The authors of (Gholami et al., 2015) generate optimal patrol strategies by modelling the relationships between the patrol unit and the criminals as a DBN.

1.2. The Process

A systems engineering process called a "spiral process" (Boehm, 1988) was followed in constructing and evaluating the causal model. Our take on the spiral process is illustrated in Figure 1 and it is called such because of its continuous nature of structuring, analysing, and synthesising, while moving ever closer to a desired solution.

According to Korb & Nicholson (2010), the spiral model for knowledge engineering with BNs (KEBN) starts off with a requirements analysis, then proceeds to the design phase, and is then implemented after which it is validated and tested. The process repeats itself until some stopping criteria is reached. This study loosely follows this process except for the implementation phase: the goal is to predict the area of possible poaching, thus it would not make sense to implement a first-draft model in the field. At this stage of the study the model has gone through three phases of requirements analysis and design, as well as two phases of validation (by domain experts).

A group of BN experts developed a first-order BN model (also called the "current perspective" model) through studying the literature and theory base (Koen et al., 2014). This current perspective BN was presented at a workshop consisting of domain experts. The current perspective model was used instead of the traditional "clean slate" approach where the experts have to come up with their own model during the workshop. Some of the main reasons for bringing a first-order BN to the workshop were to eliminate time and availability constraints. Owing to the sensitivity of the problem, experts might need to leave the workshop at a moment's notice. The workshop was

also held roughly 450 kilometre from where most of the experts work, thus travelling time and availability for that time period had to be factored in.

The current perspective model was evaluated and modified by the experts, after which subnetworks were identified that could be used to simplify future elicitation. Probabilities were elicited from experts as well as from literature to populate the model. Individual meetings with key experts who could not attend the workshop were set up and the model was finalised up to the point of general agreement.
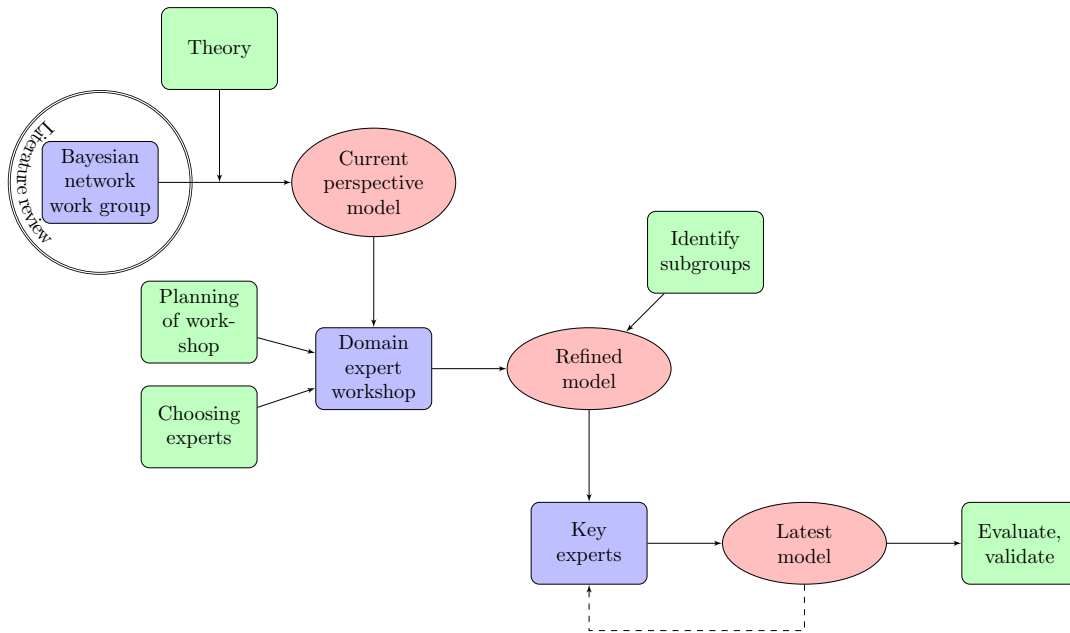
Figure 1: Spiral-type process of structuring - analysing - synthesising

This paper focuses on the last two nodes in Figure 1: the latest (expert knowledge) model and its evaluation and validation. Sensitivity analyses are performed and experts evaluate the latest model. Section 2 presents the model in its latest form, and Sections 3 - 6 describe the modelling of the different aspects of the problem. Section 7 evaluates the model by means of case-based analyses, and Section 8 concludes the paper.

2. Materials and Methods

Figure 2 depicts the problem and modelling spaces. The large oval contains a representation of everything relating to the poaching environment, and everything outside it represents that which is unknown about the rhino poaching problem. Some of these unknown elements are socio-economical variables that are important for understanding the problem in its entirety, but are not addressed in detail in this paper. The coloured nodes depict the aspects of the problem that are modelled, for example *Active_deterrents*, *Poacher_present*, *Rhino_present*, and *Poaching_event*, which are expanded on in Section 2.1. These four nodes form the main building blocks of the rhino poaching problem, and is one of the outcomes from discussions with experts. The arrows point towards the target node of the model (*Poaching_event*) which yields the probabilities that we are interested in.

The map of the KNP is divided into 1×1 kilometre cells. The same BN model is used for each cell, where evidence specific to that cell is entered into the BN. A poaching event probability is calculated using this evidence on a cell-by-cell basis.

From the Routine Activity Approach it follows that for a crime to occur (*Poaching_event* = "True"), three conditions have to be met (Cohen & Felson, 1979). There has to be (1) a victim or target (*Rhino_present* = "True"), (2) a criminal (*Poacher_present* = "True"), and (3) the absence of a capable guardian (*Active_deterrents* = "Not effective").
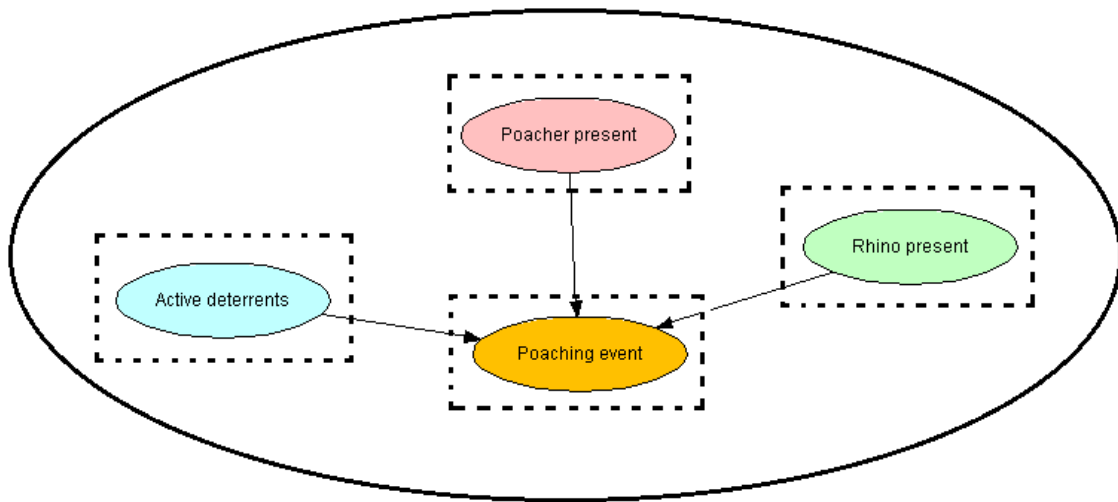

Figure 2: The big picture

7

Experts were consulted owing to the scarcity and incomplete nature of the data. The elements within the model each represent a subgroup of the complete model after consultation with experts, as is shown in Figure 3.

2.1. A Detail-Oriented Perspective

An in-depth analysis can now be made of the experts' motivation behind each subgroup and each node contained therein. The causal model presented in Figure 3 is a Bayesian network (BN) that represents the causalities in the model (Neapolitan, 2003). Nodes represent variables and arcs represent causality, and the causalities are connected to certain probabilities. Each node has a conditional probability table (CPT) describing the relationship between the parent nodes and the child nodes.

Probabilities for the BN can be supplied in a number of ways. Firstly, the root nodes (nodes that do not have parents) are populated with probabilities either from experts, the literature, or data. If we study the *Weather* node (assume that it has two states) and we know that it rains half of the time, we can say that $P(Weather = \text{"Rainy"}) = 0.5$, and $P(Weather = \text{"Sunshine"}) = 0.5$. The probabilities of non-root nodes are then calculated using Bayes' rule. To calculate the probability of a certain node's state requires only knowledge of the nodes directly influencing it. Thus, a child node is only dependent on its parent nodes. An example of this is supplied in Table 1 where the conditional probability table (CPT) for *Poaching_report* is shown. *Poaching_report* is conditioned on *Poaching_event* (the parent node) and there is a probability assigned to each combination of the states of the two variables. Owing to the sensitivity of the problem, the rest of the CPTs can unfortunately not be divulged. The probabilities for non-root nodes in a BN are calculated using Bayes' rule. The conditional probability distribution, conditioned upon available evidence, is used to perform inferences on the model.

Table 1: *Poaching_report* CPT

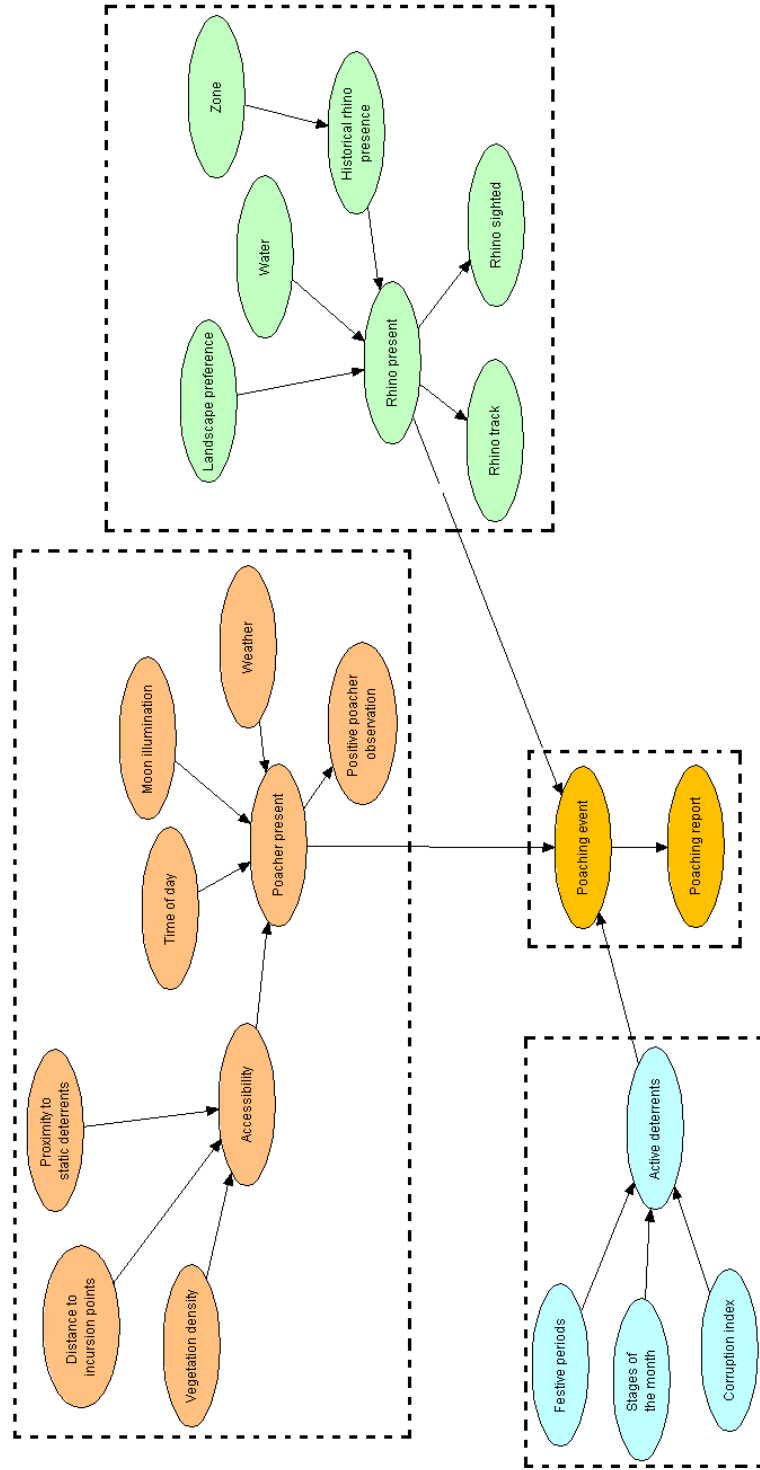| States		Probabilities
Poaching_event	*Poaching_report*	Probabilities
False	False	0.9990
True	False	0.0500
False	True	0.0010
True	True	0.9500

Figure 3: Rhino poaching model

The model is partitioned into subgroups, and each subgroup contains its own small network. Partitioning the network like this simplifies the discussions around the network, as well as obtaining probabilities for the CPTs. The network will be discussed at the hand of this partition.

The reason for using BNs is twofold: the modeller wants to understand the problem and also wishes to make predictions about it. Understanding goes hand-in-hand with prediction; if the modeller understands, then the modeller can predict. If the structure of the rhino poaching problem was fully understood the modeller could have used different methodologies such as ontologies (Guarino, 1998). Ontologies yield relationships between variables, whereas we wanted causal links. BNs provide us with the opportunity to link variables causally by using data, expert knowledge, or a combination of both. In this study only expert knowledge was used. Spatial variables were processed using the available digital maps, as well as expert knowledge and literature. The following sections describes and analyses each node in the model in Figure 3.

3. Modelling the Poacher's Presence

The *Poacher_present* subgroup contains everything we know about poacher presence. Nodes influencing a poacher's presence contain a mixture of spatial and temporal variables and is presented by the red variables in Figure 3.

3.1. Vegetation density

This node establishes where in the KNP there are areas that are very dense that could limit the poacher's access to other parts of the KNP. The states of *Vegetation_density* divide the KNP into areas that are very dense and areas that are open. Open areas are classified as any areas where there are not very dense vegetation. This node does not state which type of vegetation density is preferred by a poacher, it only states the probability of finding very dense vegetation and open areas in the KNP. A "land type" shapefile was used to compute the prior probabilities for *Vegetation_density*. Each cell of the map was classified as having either very dense vegetation or open areas by computing the area made up by each and deciding which one covers the largest area in the cell.

3.2. Distance to incursion points

Incursion points or crossing points are locations where poachers can easily gain entry to the park or exit the park. Incursion points can be holes in the fence, river crossings, or pick-up points at the side of roads.

The states of *Distance_to_incursion_points* divide the KNP into cells that are less than two and a half kilometres from an incursion point, cells that are between two and a half and five kilometres from an incursion point, and cells that are further than five kilometre from an incursion point. A shapefile containing known crossing points were used to process this root node. The distance to incursion points were calculated from the actual incursion point to the midpoint of the cell. This variable does not prioritise the probabilities according to poacher preference, it only yields the probability of being in a cell that is in the proximity of an incursion point.

3.3. Proximity to static deterrents

This node quantifies the proximity of a cell to a static deterrent, where a static deterrent can be a rest camp, lodge, ranger post, patrol hut, or gate. Static deterrents are fixed, although the location is not necessarily known by individuals other than the modeller.

It is believed that poachers prefer to stay away from these locations for fear of being detected, thus it is seen as a deterrent. Rhinos usually also avoid busy areas. This node indicates the probability of being in a certain proximity to static deterrents given a random location in the KNP. It divides the KNP into cells that are less than five kilometres from a static deterrent, cells that are between five and 10 kilometres from a static deterrent, and cells that are further than 10 kilometres from a static deterrent.

3.4. Accessibility

The *Accessibility* node can be seen as the combination of the previous three nodes and states the level of ease for a poacher to enter the park and get to the rhinos. *Accessibility* is the parent node and is influenced by *Vegetation_density*, *Distance_to_incursion_points*, and *Proximity_to_static_deterrents* (the child nodes). This node does not influence whether there is a rhino present or not, but it does influence whether or not the poacher can get to the rhinos. *Accessibility* is a hidden node as it cannot be observed directly.

3.5. Time of day

Rhino carcasses are usually detected long after poaching events. The KNP relies on necropsy results to estimate the date of death, but necropsies cannot be used to estimate the time of death when the carcass is more than a few days old. The only time this node can be populated is if a fresh carcass was found and analysed, as the necropsy will then be able to give a reasonable time of day. The states for *Time_of_day* is "Midday", "Twilight", and "Night" which were chosen to

cover an entire day, even though most of the poaching events are assumed to occur during twilight or at night.

3.6. Moon illumination percentage

It became evident from various communications that the light of the moon plays an important role in poaching events. *Moon_illumination_percentage* is a discrete variable with four states corresponding to four equally wide percentage intervals corresponding to how much light is available: $0 - 25\%, 25 - 50\%, 50 - 75\%, 75 - 100\%$.

Moon_illumination_percentage plays a role in how much can be seen at night. This node does not tell us which moon phase is preferred by poachers, it just gives the probability of *Moon_illumination_percentage* being in a certain state on any given date. A Python software package was used that takes a date as input and gives the corresponding moon illumination percentage as an output. The dates for an arbitrary calendar year were used as input to classify the dates for one year according to the four moon illumination percentage states. The number of days falling in each state were counted and divided by the number of days in the calendar year to obtain a prior probability for *Moon_illumination_percentage*.

3.7. Weather

Weather could play an important part in the poaching strategy. If it rains the rain washes away a poacher's scent so that sniffer dogs cannot follow him. In very hot and sunny conditions the scent of the poacher disappears quickly, thereby inhibiting the ability of sniffer dogs to follow his trail.

Weather is a discrete variable with four states corresponding to the four main weather types found in the KNP: "Clear", "Overcast", "Rainy", and "Windy". These four weather types were chosen to cover all the main weather conditions in the park. There are only two main seasons in the KNP, namely a wet summer and a dry winter. The dry winter season ranges from May to September/October and the wet summer season ranges from November to April. From November onwards there is high rainfall, and the summer rains peak in January/February and end in April.

The prior probabilities were extracted from literature concerning traditional weather patterns in the KNP and in Limpopo, as well as from personal experience. This was done as no accurate historical weather records were available at the time. This node does not state which weather type is preferred by poachers, but it gives the prior probability of the weather being in any one of the four states given a certain date and time.

3.8. Poacher present

The *Poacher_present* variable can be seen as the combination of the previous four variables described in this subgroup, and states whether there is a poacher present in the cell or not ("False" and "True"). All the factors that influence a poacher's presence in the park, such as *Accessibility*, *Time_of_day*, *Moon_illumination_percentage*, and *Weather*, are included. *Poacher_present* is a hidden variable as it cannot directly be observed.

3.9. Positive poacher observation

This node is conditioned on the *Poacher_present* node as the presence of a poacher in the cell increases the likelihood of observing a poacher in that cell. It has states "False" and "True", depending on reports of poacher sightings or from tracks that were observed. Data for sightings alone are scarce, but there are data for tracks, thus all these possible observation types were included in one node.

4. Modelling the Rhino's Presence

The rhino subgroup contains everything that is known concerning the presence of a rhino in a specific cell. The subgroup is illustrated by the green nodes in Figure 3.

4.1. Landscape preference

The assumption is that a specific cell will have a higher probability of containing a rhino if it contains suitable vegetation, or is at least close to a cell with suitable vegetation. White rhinos are grazers, thus they only eat grass, unlike black rhinos that are browsers and eat the leaves of low trees and shrubs.

It proved very complicated to conduct a comprehensive study of the vegetation preferences of white rhinos, thus it was decided to "keep it simple". After communicating with wildlife management and ecology experts and consulting various sources, it was decided to combine all the habitat and landscape variables into a single variable called *Landscape_preference*. The node informs us where the rhinos are and thus encapsulates the vegetation type, the grazing availability, *etcetera*.

The works of Pienaar et al. (1992), Pienaar et al. (1993), and Pienaar (1994) together with a land types shapefile containing information about the woody cover, the geology of areas, as well as the soil types, were used to process *Landscape_preference*. The cells of the KNP were classified according to these sources as being avoided, neutral, or preferred landscapes of white rhinos.

4.2. Water

White rhinos will drink water regularly in wet seasons and undertake long journeys in dry seasons to get to water (Pienaar, 1994). Owing to the importance of water, a cell with a water source will have a larger probability of having a rhino present than a cell without a water source.

The calculations for $Water$ were performed by using shapefiles containing water points and main rivers. The water point shapefile was created by merging the shapefiles of waterholes, springs, water points, water pans, and water dams.

4.3. Historical rhino presence

This node states what the distribution of rhinos previously looked like in the past in the KNP. If the map of the past four years shows a clustering of rhinos in a certain area, then it can be assumed that rhinos will be in that area this year as well. Rhino sighting data from 2011 to 2014 were used to process the $Historical\_rhino\_presence$ node. The prior probabilities were calculated according to how many rhinos, on average, frequented that cell in the past.

One of the challenges of BNs is finding the optimal set (and number) of states for a node. In this case it was decided to just use four states, as more states would cause the CPT to expand to a very large size. The eight neighbours of a cell were calculated and the total number of rhino sightings for those nine cells were added. This number was then the total number of sightings for that cluster and was used to compute the four states for classifying the cells of the map. The cells were then categorised according to this number.

4.4. Rhino present

This node can be seen as a combination of the previous three variables that influence the possible presence of a rhino. These variables are $Landscape\_preference$, $Water$, and $Historical\_rhino\_presence$. $Rhino\_present$ is a hidden binary variable with states "False" and "True" depending on the presence of rhinos in a specific cell.

4.5. Rhino track

$Rhino\_track$ is an observation variable stating whether some rhino tracking system indicates that rhinos are in a specific cell (Koen et al., 2014). $Rhino\_track$ is conditioned on $Rhino\_present$. The presence of a rhino in a specific cell increases the chances of it being tracked.

The states are "False" and "True" depending on whether or not tracked rhinos are present. There will be no tracking data if there is no rhino present, except if the tracking system tracks a different type of animal but concludes that it is a rhino.

4.6. Rhino sighted

The *Rhino_sighted* node is used to enter evidence into the network in the form of a rhino sighting in a specific cell (Koen et al., 2014). This node has two states ("False" and "True") depending on whether or not a rhino was sighted.

If a rhino is present, it will increase the probability of a rhino being seen. However, the probability of observing a rhino when there is a rhino present is quite small as rhinos are, despite their size, very adept at hiding in tall grass or under shrubs or bushes. Rhinos cannot always be observed, thus we have to rely on other information such as sightings or tracking data.

5. Modelling the Effectiveness of Deterrents

The ranger subgroup contains everything known about how effective rangers and other safeguarding measures are to protect rhinos. The subgroup is shown by the blue nodes in Figure 3.

5.1. Corruption index

According to the Routine Activity Approach (Cohen & Felson, 1979), the absence of rangers combined with the presence of both a poacher and a rhino increases the probability of a poaching event. According to experts, there can be a measure of corruption or "Ill-discipline" present in the park (Expert workshop, personal communication, September 29 - October 1, 2014) that could hinder operations. The states of *Corruption_index* are "Corrupt" and "Not corrupt". There is no data available for this variable, thus it was decided to assign a small weight to the "Corrupt" state.

5.2. Stages of the month

According to experts, poaching events appear be linked to payday (Expert workshop, personal communication, September 29 - October 1, 2014). The *Stages_of_the_month* node influences whether or not the rangers (as well as other deterrents) are effective in hindering poachers.

The four states of *Stages_of_the_month* are the week before payday, the week after payday, two weeks after payday, and the rest of the month. The probabilities were calculated by counting the number of days per month that fall in an interval, and then computing the ratio with respect to the total number of days in a month.

15

5.3. Festive periods

It appears that there could be certain poaching spikes during the year that correspond to holidays and festive periods in South Africa and Mozambique. The *Festive_periods* node influences the effectiveness of the deterrents. The states of *Festive_periods* are "False" and "True" depending on whether the specific date falls in a festive period.

Multiple calendar years were used to count how many days in a year fall in the main festive periods. It was decided to include the main festival periods as well as the week before and the week after each festive period. The ratio of festival periods to the number of days in the year was then calculated and used as the prior probabilities for *Festive_periods*.

5.4. Active deterrents

The *Active_deterrents* node tells us if deterrents are effective in deterring poachers. Active deterrents are the elements put in place to stop rhino poaching such as ranger patrols, sensors, alarms, lights, *etcetera*. The states for *Active_deterrents* are "Not effective" and "Effective". *Active_deterrents* can be seen as the combination of three abstract factors relating to the effectiveness of deterrents, namely as *Corruption_index*, *Stages_of_the_month*, and *Festive_periods*.

6. The Poaching Event

The poaching event subgroup is the key subgroup of the model and contains everything we know about whether or not a poaching event is going to take place. The subgroup is illustrated by the orange nodes in Figure 3.

6.1. Poaching event

This node forms the core of the model as *Poaching_event* is the target variable. *Poaching_event* is a binary variable that takes on the value "True" if a poaching event will take place in a specific cell, and "False" otherwise. This is a hidden variable because it cannot be observed directly.

Poaching_event can be seen as the combination of the previous three factors leading to a poaching event, namely *Active_deterrents*, *Poacher_present*, and *Rhino_present*. The ideal situation for a poaching event to occur is if a poacher is present, a rhino is present, and the active deterrents are not effective. A poaching event cannot take place if a rhino or a poacher is not present.

16

6.2. Poaching report

This node is an observed binary variable that takes on the value "True" if a poaching report is generated, and "False" otherwise. If a poaching event occurs, the probability of a poaching report being generated is high, unless someone wants to hide any record of the poaching event. If a poaching event did not occur, there is nothing to report, hence there will be no poaching report.

7. Evaluating the Expert Knowledge Model

7.1. Gaining a Better Perspective

At the time when the first-order model was built, very little usable data was available, thus expert knowledge was used. Experts add hidden variables to a model to quantify their beliefs in certain events. The challenge with this, however, is that data sometimes do not exist for these hidden variables. These same challenges occur in the rhino poaching model. Even though more usable data now exist, this particular model cannot be trained in its current form. There does not exist data for all the variables (for example $Time\_of\_day$), and for some variables there will never be any data (such as $Accessibility$). Hidden variables also play an important role in simplifying a network. Referring back to Figure 3 and the $Accessibility$ node: if this node did not exist, a total of six nodes would feed into $Poacher\_present$ ($Vegetation\_density$, $Distance\_to\_incursion\_points$, $Proximity\_to\_static\_deterrents$, $Time\_of\_day$, $Moon\_illumination$, and $Weather$), making it an ever bigger challenge to populate the conditional probability table than it is now.

The output of the BN is in the form of a heat map illustrating possible poaching hotspots. This heat map can be shown to the operations managers and they can allocate resources more effectively. An example of such a heat map is shown in Figure 4. Owing to the sensitivity of the problem, only a small area of the KNP map can be shown. The white areas correspond to areas with a high probability of poaching, and the dark colours to a very low probability. Each cell in this particular map is 1×1 kilometre.

7.2. Evaluation

If a model is parameterised with empirical data, the model is tested on data. So what do we do if our model is parameterised on expert knowledge? One way of verifying the validity of the expert model is to perform case-based evaluation. Case-based evaluation infers a set of test cases and evaluates the results using expert judgement (Korb & Nicholson, 2010). Whereas Korb &
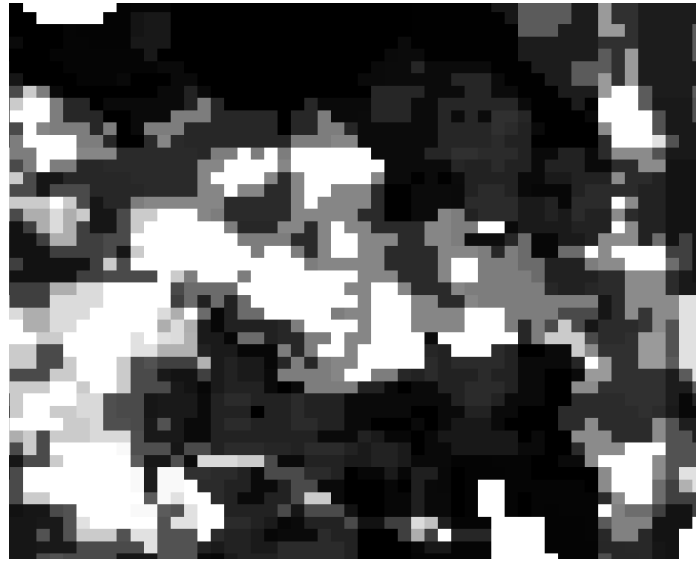
Figure 4: Zoomed in heat map: The white areas correspond to areas with a high probability of poaching, and the dark colours to a very low probability. Each cell in this particular map is 1×1 kilometre.

Nicholson (2010) suggest boundary-value analysis by choosing extreme values as input, we use the following queries for case-based evaluation: prediction, intervention and counterfactuals. These queries demand knowledge and causal stability with increasing level of details (Pearl, 2000). In this section, these three queries are performed as test cases and are discussed. The inferred results were communicated to experts at the KNP and the results are presented in Section 7.3 and discussed in Section 7.4.

Evidence nodes are nodes of which the states are observed (known with 100% probability) and are input nodes. Query nodes are nodes for which the states are unknown and are inferred by the Bayesian network inference algorithm (for example the junction tree algorithm). Query nodes are the output nodes.

7.2.1. Query: Predictions

In top-down reasoning (downward propagation), the root nodes are observed and the effects are predicted (Korb & Nicholson, 2010). For example, evidence is given for root nodes (the weather is clear, the proximity to water points is less than five kilometres, *etcetera*), and the probability of a poaching event is inferred. Figure 5 shows a screen shot of the Hugin® screen used to perform the predictive query.

In each case-based evaluation scenario the states that are observed (states fixed with a probability of 1) are shown in red. The states shown in green are those with the highest probability for each remaining node. In this specific evaluation we are interested in the probability of *Poaching_event*. The particular states of the nodes were chosen by the authors (*Vegetation_density* = "Open areas", *Time_of_day* = "Twilight", *etcetera*) as shown in Table 2. As certain states of certain nodes are fixed, the software program applies Bayes' rule to compute the probabilities for the other nodes' states, based on the states that were fixed.

The probability of a poaching event is very low because each instance of the BN corresponds to a cell in the map and, when normalised, these probabilities sum to one. A rhino poaching event is also technically a rare event, thus the probability will not be high. The important thing to note is that these are relative probabilities. Cell C_i might have a poaching event probability of 0.0348, but if cell C_{i+1} has a poaching event probability of 0.0100, then cell C_i is still three times as likely to contain a poaching event.

The combination of *Vegetation_density* = "Open areas", *Distance_to_incursion_points* = "2.5 - 5 km", and *Proximity_to_static_deterrents* = "5 - 10 km" yields a relatively low value for *Accessibility*. This low value combined with *Time_of_day* = "Twilight", *Moon_illumination_percentage* = "25 - 50%", and *Weather* = "Overcast" yields a low probability for *Poacher_present*.

In the rhino subgroup the combination of *Landscape_preference* = "Neutral", *Water* = "Close", and *Historical_rhino_presence* = "Between 13 and 26" yields a surprisingly small probability for *Rhino_present*. The probability of *Active_deterrents* = "Not effective" is quite high given that *Corruption_index* = "Corrupt", *Stages_of_the_month* = "Payday + 14", and *Festive_periods* = "False". The low probabilities of *Poacher_present* and *Rhino_present* together with the fairly high probability of *Active_deterrents* yields a very low probability for *Poaching_event*. There are no spatial variables present in the *Active_deterrents* subgroup, thus we are unsure of what the spatial impact is of *Active_deterrents* on *Poaching_event*.

7.2.2. Query: Intervention

If a poaching event does not occur, the resulting parameters change in response to interventions. This is called the prescriptive mode of reasoning and the Hugin® screenshot is shown in Figure 6. In this evaluation we are interested in the states of the other nodes when *Poaching_event* is set to "False" with a probability of 1. Here only the value of *Poaching_event* = "False" was observed, letting the software program calculate the rest of the probabilities for the remaining nodes and

19

states. The outputs of the various nodes can also be seen in Table 2.

If $Poaching\_event$ = "False" we can see that $P(Poacher\_present$ = "True") = 0.0825 and $P(Rhino\_present$ = "True") = 0.0908. Both of these values are very low and correspond to the Routine Activity Approach. The probability of $Active\_deterrents$ is more or less the same for both states, thereby saying that $Active\_deterrents$ does not really have an influence on deterring a poaching event. Once again this could be owing to the fact that there are no spatial variables present in the subgroup.
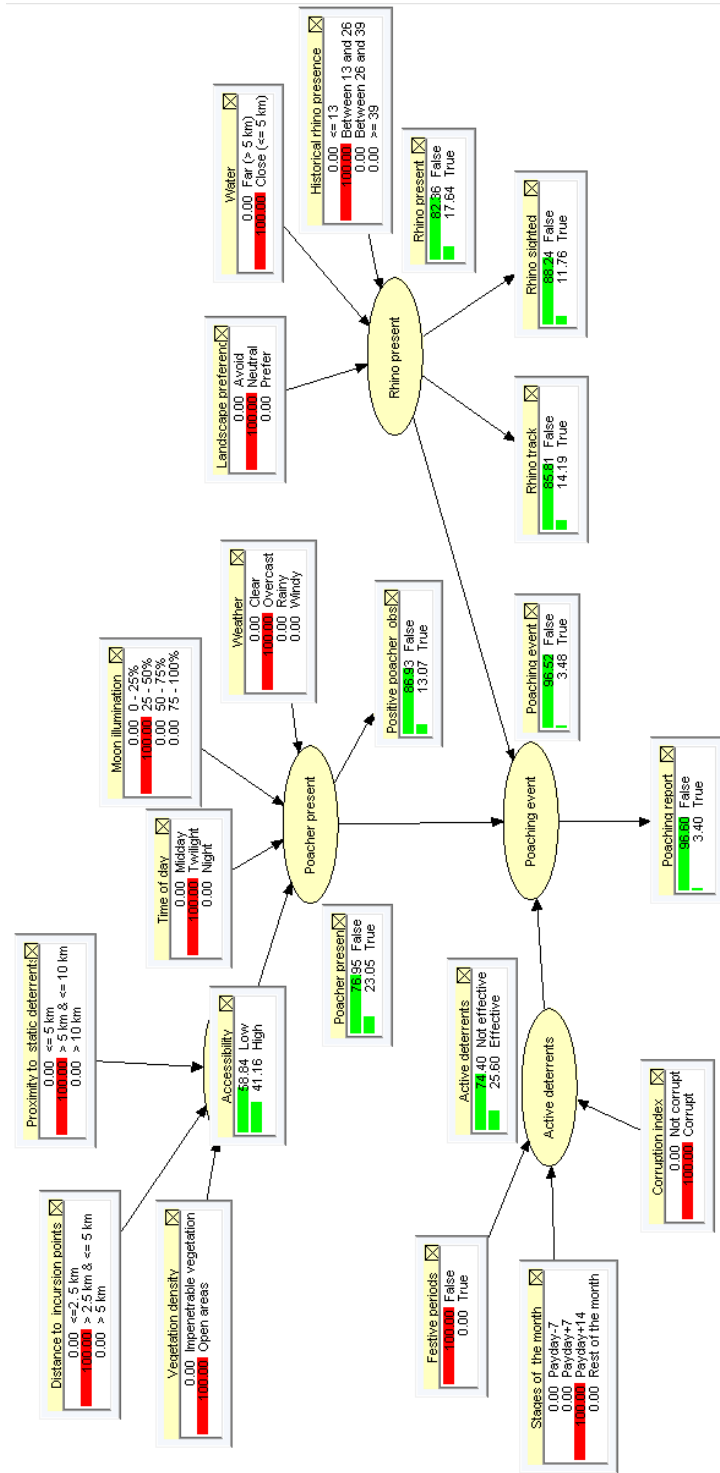
Figure 5: Query: Predictive mode (downward propagation). A state is automatically coloured red in Hugin® when the user fixes its value (evidence nodes). All the other states are automatically coloured in green. Here the target variable is *Poaching_event*. The particular states of the other nodes were chosen by the authors.

Table 2: Input versus output states for queries

| Node name	Prediction		Intervention		Counterfactual	
	Input	Output	Input	Output	Input	Output
Vegetation_density	Open areas			Open areas		Open areas
Distance_to_incursion_points	2.5-5 km			> 5 km		> 5 km
Proximity_to_static_deterrents	5-10 km			5-10 km		5-10 km
Accessibility		Low		Low		Low
Time_of_day	Twilight			Midday		Twilight
Moon_illumination_percentage	25-50%			0-25%		75-100%
Weather	Overcast			Clear		Clear
Poacher_present		False		False		True
Positive_poacher_observation		False		False		False
Landscape_preference	Neutral			Neutral		Prefer
Water	≤ 5 km			≤ 5 km		≤ 5 km
Historical_rhino_presence	13-25			≤ 13		≤ 13
Rhino_present		False		False		True
Rhino_track		False		False		True
Rhino_sighted		False		False		False
Festive_periods	False			False		False
Stages_of_the_month	Payday + 14			Rest of the month		Rest of the month
Corruption_index	Corrupt			Not corrupt		Not corrupt
Active_deterrents		Not effective		Effective		Not effective
Poaching_event		False	False		True	
Poaching_report		False		False		True
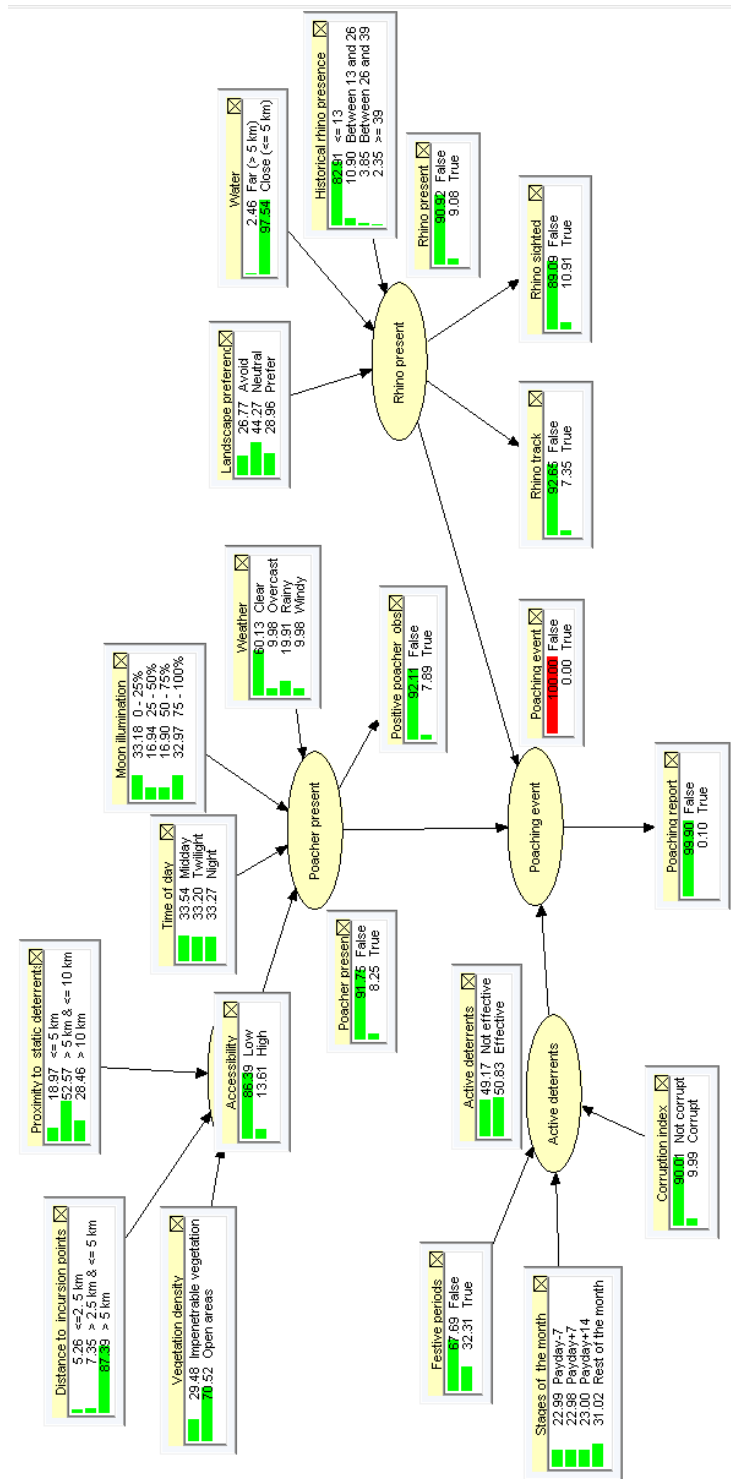
Figure 6: Query: Intervention (prescriptive mode of reasoning). A state is automatically coloured red in Hugin® when the user fixes its value (evidence nodes). All the other states are automatically coloured in green. Here the value of *Poaching_event* is fixed as "False" and the states of the other nodes are inferred.

23

7.2.3. Query: Counterfactual

Counterfactual queries pose the most demanding and distinctive characteristics of BNs. In counterfactual reasoning, the leaf nodes are observed and the causes are inferred. A "counterfactual" query analyses the causes of the (usually undesired) event and an "intervention" query analyses ways to stop the event. For example, if it is known that a poaching event did occur, available evidence can be filled in and the rest of the states that led to the poaching event can be inferred. The system is "diagnosed" as to what caused the outcome (in this case the poaching event). Figure 7 shows the screenshot of the Hugin$^{®}$ screen used to perform the counterfactual query. In this evaluation we are interested in the states of the other nodes when *Poaching_event* is set to "True" with a probability of 1. Here only the value of *Poaching_event* = "True" was fixed, letting the software program calculate the rest of the probabilities for the remaining nodes and states. The outputs of the various nodes can also be seen in Table 2.

If *Poaching_event* = "True" we can see that $P(Poacher\_present =$ "True"$) = 1$ and $P(Rhino\_present =$ "True"$) = 1$, which is in line with the Routine Activity Approach. The probability of *Active_deterrents* = "Not effective" is reasonably high, but again a conclusion cannot really be drawn as there are no spatial variables present.

If we look at the poacher subgroup it seems that *Accessibility* need not be high, but that the vegetation should be open, the cell should be far from incursion points, and at least five kilometre from static deterrents. Twilight and night seem to be preferred, as is a high moon illumination percentage. As far as weather is concerned, it seems that clear and rainy conditions are preferred. In the rhino subgroup the landscape preference should at least be neutral, the cell should be close to water, and it does not seem to matter if there were rhinos sighted in previous years.
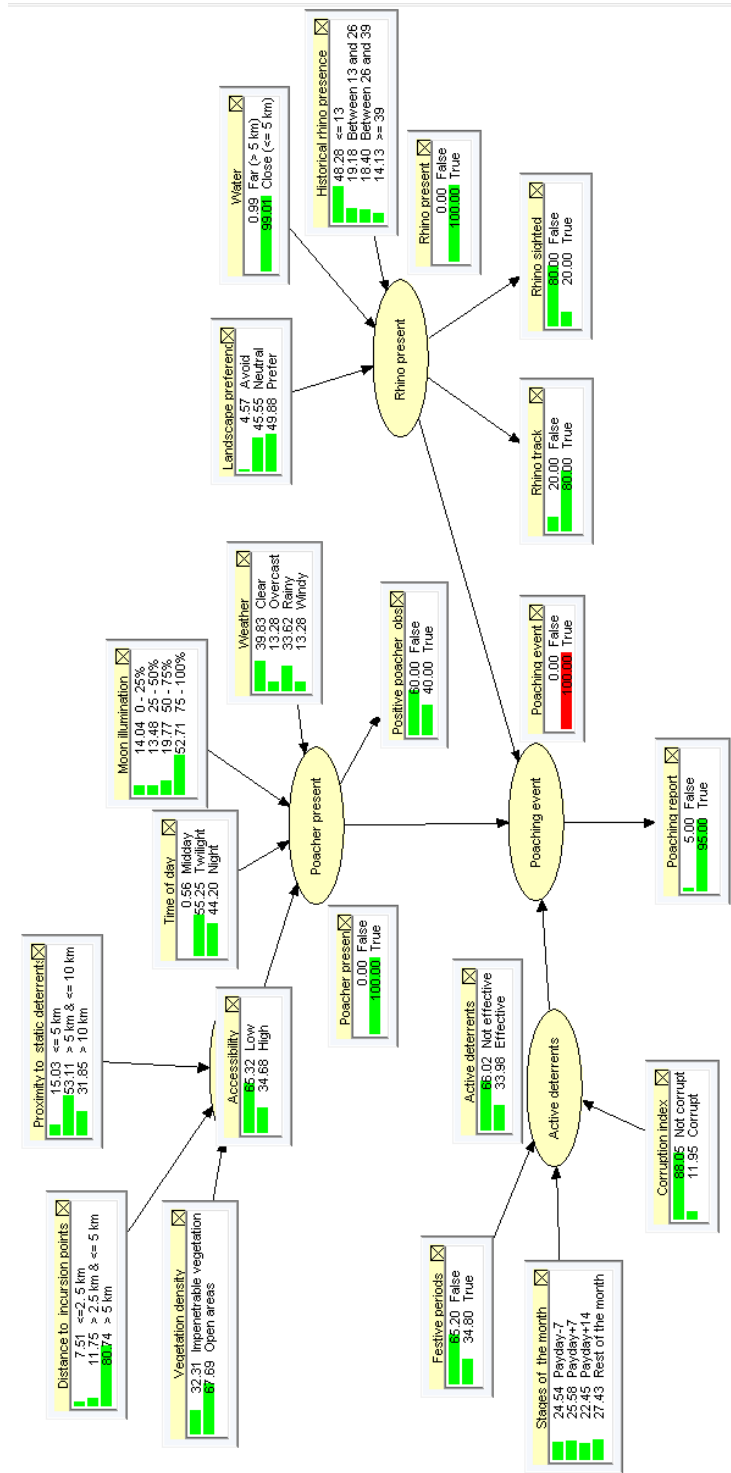
Figure 7: Query: Counterfactual (analysing the causes of the event). A state is automatically coloured red in Hugin® when the user fixes its value (evidence nodes). All the other states are automatically coloured in green. Here the value of *Poaching_event* is fixed as "True" and the states of the other nodes are inferred.

7.3. Results

Distinct from the original workshop where the model was developed, the three query methods were shown to independent experts in the KNP and they were asked to evaluate the model in terms of the outcomes of the various analyses as well as the model structure itself. The experts consisted of persons working in the field and the operations centre in the KNP. They responded positively to the model and stated that it was something that could help them make decisions about the problem. This is one of the goals we wanted to achieve with this study: to create a model that can be used to facilitate discussions and encourage new approaches.

On the whole the experts agreed with the model architecture, but suggested a few minor adjustments to the model. These adjustments will be discussed in the following section. According to the experts, the model, as it is, reflects the situation picture roughly a year ago. The variables remain the same, but the important states for a poaching event to occur has changed somewhat. Changes occur rapidly and the expert model cannot be updated unless we are kept in the loop about the changes. Owing to the sensitivity of the problem all the details cannot be divulged, but some of the adjustments and improvements are described below.

7.4. Discussion

The experts evaluated the model on a node-by-node basis. According to them, the density of the vegetation does not matter anymore for a poaching attack; the poachers will go wherever there are rhinos, thus open areas and dense vegetation carry an equal weighting, whereas a specific vegetation density seemed to be preferred two years ago.

The preferred time of day for poaching events seems to be just before twilight, which remains the same from the previous workshop. The experts suggested that the states of $Time\_of\_day$ be changed from "Midday", "Twilight", and "Night" to "First light", "Midday", "Last light", and "Night". The reason for this is that "Twilight" refers to both the time just before daybreak and the time just before nightfall. Both of these periods can be potential poaching times, thus it is important to distinguish between morning and evening to allocate resources to an area. It was also confirmed that poachers use bad weather to their advantage. In rainy conditions, normal people will seek shelter from the rain, but the poachers capitalise on that and use the rain to make their escape.

The current approach to the spatial modelling of poaching is to note where a rhino carcass has fallen instead of noting where the animals were shot. In some instances the rhino will drop

immediately, and in other cases it might travel a distance before dying. A consideration for future iterations of the model might be to include this by means of using cells adjacent to those that the rhino carcass was found in the calculations.

According to the experts there should be an extra node feeding into *Poacher_present* encapsulating the experience of the poacher. Inexperienced poachers will have a different strategy compared to experienced poachers. The one might utilise twilight and full moon, and the other might utilise the middle of the day and a dark moon at night. A node such as this might be useful in profiling the poachers, but utilising it in this model will need data, or at least viable expert knowledge.

The general consensus amongst the experts is that "corruption" is not the correct word to use in the *Corruption_index* variable. The experts stated that there is a distinction to be made between corruption and ill-discipline. A better name for this node would be *Operational_security*. Often a security breach occurs when sensitive information is leaked, which could be accidentally or intentionally. A ranger might be on a mission out in the field and phone his wife even though cellphones are not allowed in the field on that mission. This is not necessarily corruption or malice, but more likely ill-discipline.

The *Active_deterrents* subgroup in its current form is not spatially specific, thus it does not currently contribute to the overall model. This is especially true if it will not be possible to obtain spatial data for measures such as ill-discipline. If data for the patrol movements of KNP rangers are not available, it will make the spatial quantification of the active deterrent subgroup difficult.

According to the experts, the stages of the month do not play an important role in poaching anymore, which corresponds to the outcome of most of the query analyses. The rangers are dedicated and work long hours, thus reducing the time for taking leave or leaking secrets in a social situation.

The experts concur that they lose communication with their staff over Christmas and Easter and that they are usually understaffed during those periods. According to the experts, festive periods itself are not really a problem, but if they do not manage the possible poaching events during a festive period, then the months following that festive period will be hit harder than normal.

8. Conclusion

This paper presents the development of a Bayesian network to model the rhino poaching problem. Owing to the initial lack of good-quality data, expert knowledge was used to refine the structure and

populate the conditional probability tables. Expert models may also be useful in countries or game parks with limited resources for collecting carcass data. Case-based evaluation was performed to test whether or not the model coincides with logic and the knowledge of the experts. The model was shown to experts in the KNP and the outcomes of the prediction, intervention, and counterfactual queries correspond mostly to the experts' intuition. Minor changes were suggested to some of the states of the nodes and some changes that were suggested concern the weighting given to specific states. Certain states were considered more important than others a year or two years ago, but after this expert evaluation, it came to light that different states of the same variables are now considered more important. Some variables that were initially considered to be important are also not as important any more in retrospect, thus underlining the importance of updating the model frequently. This is an interesting result, as this reinforces our belief that the system is adaptive.

The conclusion that can be drawn is that this model is a tool that can be used for management decisions as well as facilitating discussions and bringing new insights to light. Structuring the problem in a causal manner led to a framework that describes the problem in a way that reveals to us what kind of data is needed in order to understand the problem better, and what the relationships between the different datasets are. This is a contribution to the field of rhino poaching and resource plundering as currently similar reasoning tools do not exist.

References

Biggs, D., Courchamp, F., Martin, R., & Possingham, H. P. (2013). Legal Trade of Africa's Rhino Horns. *Science*, *339*, 1038–1039.

Boehm, B. W. (1988). A Spiral Model of Software Development and Enhancement. *IEEE Computer*, *21*, 61–72. doi:`10.1109/2.59`.

Burn, R. W., Underwood, F. M., & Blanc, J. (2011). Global trends and factors associated with the illegal killing of elephants: A hierarchical bayesian analysis of carcass encounter data. *PLoS ONE*, *6*. doi:`10.1371/journal.pone.0024165`.

Cohen, L. E., & Felson, M. (1979). Social Change and Crime Rate Trends: A Routine Activity Approach. *American Sociological Review*, *44*, 588–608.

Critchlow, R., Plumptre, A. J., Driciru, M., Rwetsiba, A., Stokes, E. J., Tumwesigye, C., Wanyama,

F., & Beale, C. M. (2015). Spatiotemporal Trends of Illegal Activities from Ranger-Collected Data in a Ugandan National Park. *Conservation Biology*, *29*, 1458–1470. doi:`10.1111/cobi.12538`.

Duffy, R., & St John, F. A. V. (2013). *Poverty, Poaching and Trafficking: What are the Links?*. Technical Report June. doi:`http://dx.doi.org/10.12774/eod_hd059.jun2013.duffy`.

Emslie, R. H., & Brooks, M. (1999). *African Rhino: Status Survey and Conservation Action Plan*. URL: `https://portals.iucn.org/library/efiles/edocs/1999-049.pdf`.

Ferreira, S. M., Greaver, C., Knight, G. A., Knight, M. H., Smit, I. P. J., & Pienaar, D. (2015). Disruption of rhino demography by poachers may lead to population declines in Kruger National Park, South Africa. *PLoS ONE*, *10*, 1–18. doi:`10.1371/journal.pone.0127783`.

Ferreira, S. M., Pfab, M., & Knight, M. (2014). Management strategies to curb rhino poaching: Alternative options using a cost-benefit approach. *South African Journal of Science*, *110*, 1–8. doi:`10.1590/sajs.2014/20120055`.

Gholami, S., Zhang, C., Sinha, A., & Tambe, M. (2015). An extensive study of Dynamic Bayesian Network for patrol allocation against adaptive opportunistic criminals.

Guarino, N. (1998). *Formal Ontology in Information Systems: Proceedings of the First International Conference (FOIS'98), June 6-8, Trento, Italy*. IOS Press.

Haas, T. C., & Ferreira, S. M. (2015). Conservation Risks: When Will Rhinos be Extinct? *IEEE Transactions on Cybernetics*, (pp. 1–14). doi:`10.1109/TCYB.2015.2470520`.

Kahler, J. S., & Gore, M. L. (2012). Beyond the Cooking Pot and Pocket Book: Factors Influencing Noncompliance With Wildlife Poaching Rules. *International Journal of Comparative and Applied Criminal Justice*, *36*, 103–120. URL: `http://www.tandfonline.com/doi/abs/10.1080/01924036.2012.669913`. doi:`10.1080/01924036.2012.669913`.

Koen, H., De Villiers, J. P., Pavlin, G., De Waal, A., De Oude, P., & Mignet, F. (2014). A Framework for Inferring Predictive Distributions of Rhino Poaching events Through Causal Modelling. In *17th International Conference on Information Fusion (FUSION)* (pp. 1–7). IEEE Explore.

Korb, K. B., & Nicholson, A. E. (2010). *Bayesian Artificial Intelligence, Second Edition*. (2nd ed.). Florida, USA: CRC Press, Inc. doi:`10.1198/tech.2005.s836`.

Neapolitan, R. E. (2003). *Learning Bayesian Networks*. New Jersey, USA: Prentice-Hall, Inc.

Park, N., Serra, E., Snitch, T., & Subrahmanian, V. S. (2015). APE: A Data-Driven, Behavioral Model Based Anti-Poaching Engine. *IEEE Transactions on Computational Social Systems*, *2*, 15–37. doi:`http://dx.doi.org/10.1109/TCSS.2016.2517452`.

Pearl, J. (2000). *Causality: Models, Reasoning and Inference*. Cambridge University Press.

Pienaar, D. (1994). Habitat Preference of the White Rhino in the Kruger National Park. In *Symposium on Rhinos as Game Ranch Animals* (pp. 59–64).

Pienaar, D., Bothma, J. d. P., & Theron, G. (1992). Landscape Preferences of the White Rhinoceros in the Southern Kruger National Park. *Koedoe*, *35*, 1–7.

Pienaar, D., Bothma, J. d. P., & Theron, G. K. (1993). Landscape Preference of the White Rhinoceros in the Central and Northern Kruger National Park. *Koedoe*, *36*, 79–85.