PREDICTIVE POLICING IN AN ENDANGERED SPECIES CONTEXT: COMBATING RHINO POACHING IN THE KRUGER NATIONAL PARK

by

Hildegarde Suzanne Koen

Submitted in partial fulfillment of the requirements for the degree
Philosophiae Doctor (Electronics)

in the

Department of Electrical, Electronic and Computer Engineering
Faculty of Engineering, Built Environment and Information Technology

UNIVERSITY OF PRETORIA

February 2017

SUMMARY

PREDICTIVE POLICING IN AN ENDANGERED SPECIES CONTEXT: COMBATING RHINO POACHING IN THE KRUGER NATIONAL PARK

by

Hildegarde Suzanne Koen

Promoter(s):	Dr. JP de Villiers & Dr. Henk Roodt
Department:	Electrical, Electronic and Computer Engineering
University:	University of Pretoria
Degree:	Philosophiae Doctor (Electronics)
Keywords:	Rhino poaching, predictive modelling, shared awareness, causal networks

Approximately three rhinos are poached daily in South Africa. Rhino poaching is a serious problem that affects not only the rhino population of South Africa, but also the rhino population of the world. South Africa has the largest rhino population and of those rhinos the largest number can be found in the Kruger National Park (KNP). The KNP has been hit the hardest by the poaching epidemic, losing 1,175 rhinos in 2015 alone. Two big challenges are the size of the park and the unknown locations of both the poachers and new poaching events. The KNP is the size of a small country and there are simply not enough rangers to patrol this area effectively. A costly solution would be to employ more rangers, but a proposed alternative is to reduce the search space and thus ensure that the rangers are allocated to the high risk areas first. A mathematical model was developed in the form of a Bayesian network (BN) to infer the most important factors contributing to poaching events and to model the rhino poaching problem. This model can be used to predict the area in which a future poaching attack could take place and thereby reduce the search area of rangers. The model also serves as a vehicle to enhance the understanding of the problem and encourage reasoning and discussion amongst decision makers. The map of the KNP is divided into cells and each cell is given a poaching probability, based on the outcome of the BN. This probability map forms a heat map that can be shown to the control centre and rangers can then be sent to the respective hotspots on the map. This is a proactive approach,

which is in stark contrast to the numerous reactive approaches attempted thus far. This is the first BN modelling approach to the rhino poaching problem, and it is also the first BN application to wildlife crime. Previous applications of BN have only gone so far as environmental modelling, but not wildlife crimes. In this study the rhino poaching problem was shifted from a complex, ill-structured space to a position where researchers can begin to address the underlying problems by using a causal model as the vehicle for understanding the complex interplay between the factors affecting poaching events.

OPSOMMING

VOORSPELLENDE POLISIËRING IN ’N BEDREIGDE SPESIE KONTEKS: BESTRYDING VAN RENOSTERSTROPERY IN DIE KRUGER WILDTUIN

deur

Hildegarde Suzanne Koen

Promotor(s):	Dr. JP de Villiers & Dr. Henk Roodt
Departement:	Elektriese, Elektroniese en Rekenaar-Ingenieurswese
Universiteit:	Universiteit van Pretoria
Graad:	Philosophiae Doctor (Elektronies)
Sleutelwoorde:	Renosterstroping, bedoelingskatting, gedeelde bewustheid, kausale netwerke

Ongeveer drie renosters word daagliks in Suid-Afrika gestroop. Renosterstroping is ’n ernstige probleem wat nie net die renosterbevolking van Suid-Afrika raak nie, maar ook die res van die wêreld. Suid-Afrika het die grootste renoster bevolking in die wêreld, en die grootste getal van dié renosters word in die Kruger Nasionale Park (KNP) aangetref. Die KNP word die ergste geraak deur die stropings epidemie en 1,175 renosters is in 2015 gestroop. Twee groot uitdagings is die grootte van die park, asook die onbekende posisies van beide die stropers en die nuwe stropingsaanvalle. Die KNP is die grootte van ’n klein land en daar is eenvoudig nie genoeg veldwagters om hierdie area effektief te patrolleer nie. ’n Duur oplossing sou wees om meer veldwagters aan te stel, maar ’n alternatief is om die soekarea van die veldwagters te verklein en sodoende te verseker dat die veldwagters die hoë-risiko areas eerste, en meer gereeld, patrolleer. ’n Wiskundige model in die vorm van ’n Bayesiese netwerk (BN) is ontwikkel om die belangrikste faktore te bepaal wat bydra tot stropingsaanvalle en uiteindelik die probleem te modelleer. Hierdie model kan gebruik word om die area te voorspel waar ’n stropingsaanval moontlik kan plaasvind en die soekarea van die veldwagters te verminder. Dit dien ook as ’n kanaal om die begrip van die probleem te verbeter en redenasie en bespreking onder besluitnemers aan te moedig. Die kaart van die KNP word in selle verdeel en aan elke sel word ’n stropingswaarskynlikheid toegeken gebaseer op die uitkoms van die

BN. Hierdie waarskynlikheidskaart vorm ’n “hittekaart” wat aan die kontrolesentrum gewys kan word, en veldwagters kan dan na die onderskeie responskolle op die kaart gestuur word. Hierdie pro-aktiewe benadering is in teenstelling met huidige reaktiewe benaderings. Hierdie is die eerste BN modellering benadering tot die renosterstropingsprobleem, en dit is ook die eerste BN toepassing tot natuurlewe-misdaad. Vorige toepassings van BNs het omgewingsmodellering aangespreek, maar nie natuurlewe-misdade nie. In hierdie studie word aangetoon hoe die renosterstropings probleem geskuif is vanaf ’n komplekse, swak gestruktureerde probleemruimte na ’n omgewing waar navorsers kan begin om die onderliggende probleme aan te spreek deur gebruik te maak van ’n kausale model as die voertuig van begrip om die komplekse wisselwerking tussen faktore wat ’n stropingsaanval beïnvloed, te verstaan.

TABLE OF CONTENTS

CHAPTER 1 INTRODUCTION

1.1 INTRODUCTION

In 2012 you could not turn on the television or radio, or open a newspaper, without hearing or reading articles about rhinoceros (rhino) poaching and what the public can do to protect and save the rhinos. Pictures were painted of rhinos bordering on extinction, and awareness was created for saving these creatures. Since then the media attention has lessened somewhat and the public has lost interest or found new causes to support. The reality, however, is that the poachers have not lost interest and rhinos are still being slaughtered.

Rhino poaching is a crime that has increased at an alarming rate in recent years. It affects the rhino population of South Africa and by extension, the world. The next few sections discuss the significance of the problem as well as the causes. Current approaches as well as the approach followed in this thesis is also discussed.

1.2 PROBLEM STATEMENT

The problem we are faced with is the poaching of rhinos at a rate that increases rapidly each year, possibly up until the point of extinction. The number of rhino poaching attacks in South Africa has increased significantly since 2008 as can be seen in Figure 1.1. In 2008, a total of 83 rhinos were poached in South Africa and that increased yearly to a total of 1,215 rhinos poached in 2014 [1] and 1,175 rhinos poached in 2015. Figure 1.2 illustrates the poaching figures for 2014. Since 2015 the poaching statistics are not disclosed publicly anymore, thus a graph for 2015's poaching statistics is unavailable. There has been a decrease in the number of poachings in 2015, but this decrease is slight.

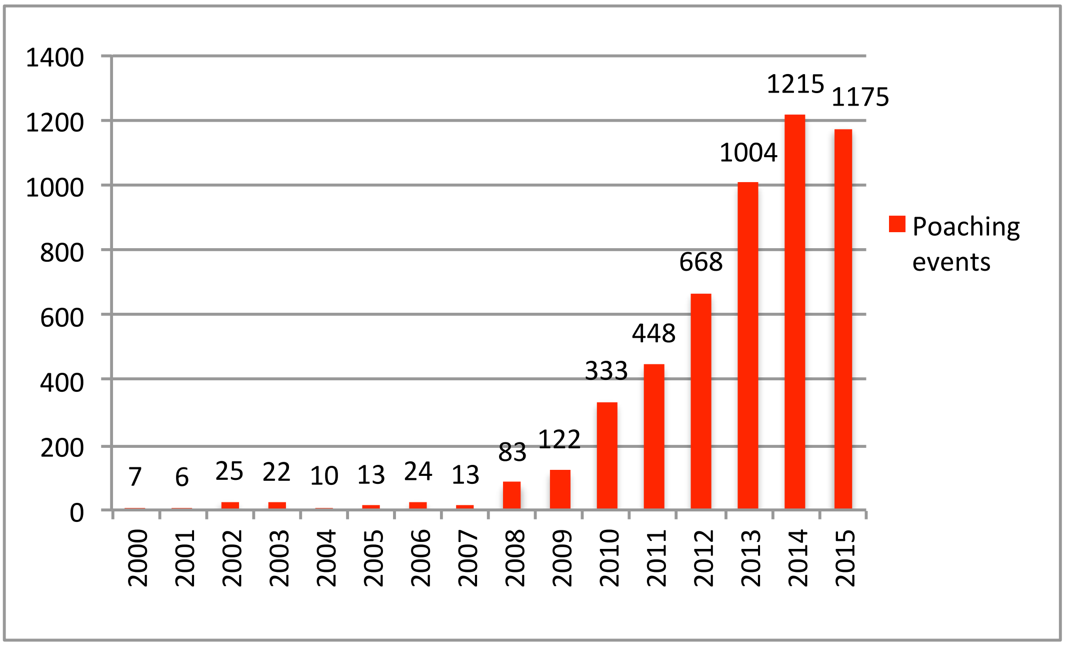

Figure 1.1. Poaching figures for 2000-2015

The Kruger National Park (KNP) has the largest population of white rhinos in the world, and the second largest population of black rhinos in the world [2]. The KNP has also been hit hardest by the poaching attacks (as can be seen from Figure 1.2). During 2012 and 2013 the KNP suffered 60-64% of the poaching attacks, while this figure has increased to 68% in 2014 [1].

One of the main reasons for the onslaught on the KNP's rhino population is the fact that the KNP is located on the border between South Africa, Zimbabwe, and Mozambique. Zimbabwe and Mozambique are known to have poaching syndicates that are involved in rhino horn trading [3].

The Great Limpopo Transfrontier Park was established over a decade ago with the idea of lifting the fences and connecting all the large game reserves in these three bordering countries, thereby creating a single immense national park that spans Limpopo (a province in the northern part of South Africa that is home to most of the KNP's geographical region), Zimbabwe, and Mozambique [4, 5]. There has been a lot of difficulty in establishing this park, as the countries involved cannot seem to agree on a suitable course of action. Certain fences have already been lifted, granting both the animals and the poachers access across borders. However, owing to the political boundaries that still exist, there are

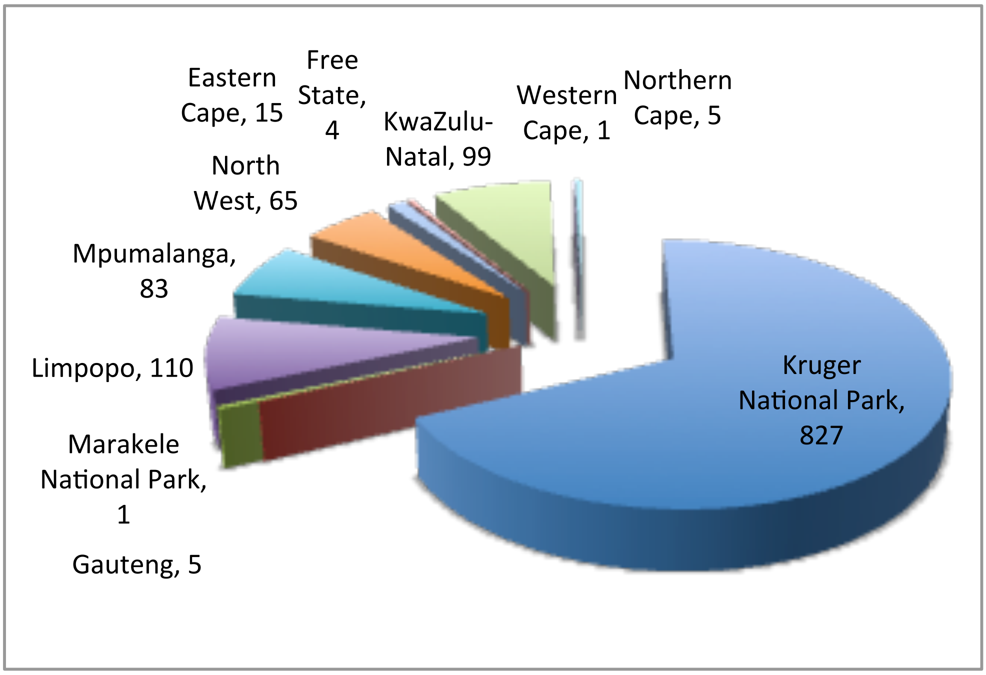

Figure 1.2. Poaching figures for 2014 according to area

still many areas that the rangers may not cross. A poacher can easily poach a rhino in South Africa and cross one of these "lifted" borders and escape to Mozambique, while the South African rangers are bound by law and cannot cross those same borders to capture the poachers.

Another reason for the poaching epidemic is the sheer size of the KNP. The KNP roughly spans 19,485 km^2 [6]. It is roughly 360 kilometre on its longest side and 65 kilometre at its widest, and is slightly smaller than Belgium. The size of the KNP makes it easy for poachers to hide amongst the vegetation and escape.

1.2.1 The problem

These poaching operations are illegal and dangerous, for both the rhinos and the people in the parks. Rhinos are poached for their horns. It is, however, incorrect to call it a "horn", since it is not made of bone, but mainly of keratin [7], the same substance human teeth, hair, and nails consist of. However, for the rest of this thesis it will be referred to as rhino "horn" to avoid confusion.

The exact reasons for the poaching epidemic are still largely unknown, but the main drivers behind it seem to be greed and poverty. Cultural beliefs and a lack of education also play big roles in rhino poaching. In many Eastern countries, rhino horn is believed to have various medicinal benefits such as curing illnesses and ailments such as arthritis [2, 8, 9], although numerous scientific studies have proven this to be untrue [10]. More recently rhino horns have also been advertised as a cure for cancer in countries such as China and Vietnam. The growing upper-middle class in these Eastern countries are also using rhino horn as a recreational drug, while many use it as a hangover cure from excessive consumption of alcohol [8, 10].

Apart from the believed medicinal benefits, rhino horn is also a status symbol. Individuals who can afford rhino horn have to be rich and have a high social standing [11]. Rhino horn is seen as a commodity, and exporting it has become a lucrative business [8]. Eloff [12] reported that in 2012 poachers received at least R81,000 per rhino horn. According to the Environmental Crime Investigation Unit poachers received a staggering R125,000 per kilogram of rhino horn in 2014 (pers comm).

Another reason for the high poaching numbers is that in many countries wildlife crime (of which rhino poaching is an example) is not really seen as a serious crime [13]. This notion is changing rapidly in South Africa where rhino poachers have recently started receiving heavy sentences. In 2014, a poacher was convicted of three counts of rhino poaching and sentenced to 77 years in prison [14]. Unfortunately, the same cannot be said for other countries, such as Mozambique [3]. In Mozambique, rhino poaching is not seen as a serious crime, but carrying an illegal firearm, however, is (Personal communication, 12 December 2014). The law regulating illegal firearms is currently used to catch rhino poachers, but the law is changing to include the crime of poaching. Currently, poachers can escape to Mozambique without facing any serious charges.

Poaching operations are usually conducted by poaching syndicates. The so-called “kingpin” employs syndicate members to order the kill and obtain the rhino horn. The actual poachers are referred to as “foot soldiers” and are usually poor men from nearby settlements [12]. There is a shortage of work and economic opportunities in these settlements, thus poaching is a lucrative option. A single poaching job can elevate an impoverished foot soldier to that of middle class in his settlement. If a poacher is caught, his punishment is in accordance with what he was caught with. The worst-case scenario would be if he was caught with a rhino horn in his possession, or with a weapon such as a panga (an African variant of the machete) or a rifle. If a poacher is caught empty-handed, he is questioned and jailed for the night where he receives a place to sleep and a meal.

The poacher is arrested and then given a fine before being released. The only time the police really have a case against the poacher is if they either catch him killing a rhino or cutting off a horn, or if he has the horn in his possession, by which time it is already too late. The goal is to prevent poaching events, and not only to react once the crime has been committed. Incarcerating poachers is a short-term solution, as it unfortunately does not solve the problem. For every poacher incarcerated, there is a handful more to take his place. When comparing the risks to the rewards, it is clear why there are so many poached rhinos in South Africa.

1.2.2 The facets of rhino poaching

According to Biggs [15], the remaining rhino population in Africa might become extinct in the next 17 years if the poaching trends of the past few years continue to increase. Duffy [16] is in agreement and states that the current rate at which rhinos are poached is less than the birthing rate, so the rhino population is still increasing. However, if current trends continue, the tipping point (where the poaching rate exceeds the birthing rate) could be reached very soon. At this point the rhino population will decline until extinction. Although there has been an increase in the number of incursions, there has been a slight decline in poaching attacks during 2015. This is due to efficient ranger activities, but unfortunately, the poaching rate remains roughly three rhinos per day.

Rhinos can be bred in private game reserves, but by doing so the gene pool could be reduced. After several decades, the gene pool might be reduced to an inferior selection. Rhino poaching is not only detrimental to the rhino population, but it is a vicious criminal circle. The ends of justice are defeated time and time again and it spills into other socio-economic areas. It has been reported that the revenue created from rhino poaching is used to fund other illegal operations such as smuggling, arms trading, and human trafficking [17].

Rhinos are also not the only animals being poached: abalone, tigers, elephants, and lions are also poached. Action needs to be taken quickly to fight these illegal operations. Stopping rhino poaching might stop the further overflow into these environmental crimes.

1.2.3 Previous and current approaches to the problem

Other countries have performed work on the subject of *predictive policing* in the context of reasoning humans (criminals) versus reasoning humans (victims). Rhinos cannot reason in the same way humans can, making the problem more challenging. Instead of reasoning humans versus reasoning humans, we now have reasoning humans (the poachers) versus rhinos. Rhinos cannot learn how to fend off poachers, or how to protect themselves; the only party that "learns", is the poachers. The game rangers need to reason for the rhinos, by proxy. Another challenge is the problem of corruption.

Various approaches have been attempted to combat the rhino poaching problem, ranging from deploying the military in the park [2], to injecting rhino horn with poison [18], to spending large sums of money on technology [8]. Very few studies use a proactive approach towards wildlife crime, and even fewer studies use mathematical prediction models to address the problem. The next chapter presents the available literature.

The military has been deployed in the KNP since 2012 [2, 19] and although they have had success in apprehending the poachers, they do not seem to have had much success in preventing the poaching attacks. According to Muntifering [20] the conservation policy for the protection of rhinos worldwide thus far has been one of increased investment in military presence and military-style enforcement. These enforcement techniques are based on simplistic models of rational crime and economics. A poacher, however, does not weigh the pros and cons of crime before poaching a rhino. According to the authors, for solutions to be effective, "...a context-specific, stakeholder-driven mix of top-down and bottom-up mechanisms grounded in theory that represents human behaviour more realistically" [20] will be needed. The authors make the case for community-based strategies and go on to state that they agree with law enforcement being a critical measure to curb poaching, but that it is not the only measure that should be used.

Certain game reserves have injected rhino horns with a chemical substance that, if consumed, will make the person extremely ill [2]. They have had some success with this as game reserves noticed a sharp decline in poaching attacks immediately after broadcasting that they have injected a substance into their rhino's horns, making them unsuitable for human consumption. The concern is that the poison might be deadly to other animals that might feed off rhino carcasses. A recent study by Ferreira [21] however debunked the claims made that the dye curbed rhino poaching.

Large sums of money have been spent on developing new technology to aid in the apprehension of poachers such as radars, sensors, and drones [2, 8]. Decision makers within the KNP have also been bombarded with an array of technological advances, each promising to be better than the previous one. This is all good and well, but night-vision goggles, for example, will not help you if you do not know where to look!

Game reserves and national parks have begun to realise the importance of having robust infrastructure and sound models before implementing all the other aspects mentioned above. A solid framework will yield the ability to slot new technologies into the system without configuring everything from scratch. This framework includes strategies such as whole-of-government approaches [22] and transdisciplinary approaches [23], and paves the way for other systems such as predictive modelling [24, 25].

Individuals seem to be in two camps about whether the legalisation of rhino horn will be a good alternative to the current situation. South Africa considered legalising rhino horn locally, but that was rejected in early 2016. Many individuals think that legalising rhino horn will not work due to the fact that the major problem is internationally, and not locally. Other individuals seem to think that the legalisation of rhino horn will not work due to the demand becoming larger. There will never be enough rhino horn to satisfy the worldwide demand, and rhinos will continue to be poached. “Cultivated” rhino horn will also not be seen as exclusive as “free-range” rhino horn, therewith driving the demand for rhino horn even further.

Another solution to rhino safeguarding is to employ more rangers. This solution will likely work if there are enough rangers to cover the whole game reserve, and yes, more jobs will be created, but the government cannot afford this. Employing enough rangers to effectively safeguard every single rhino in South Africa would have cost the government about R3.3 billion per annum in salaries in 2011 already [26].

1.2.4 Why a new approach is needed

Extra fences are currently erected, sniffer dogs are employed, and security is tightened across all game reserves. Still there are more than three rhinos poached every day. A total of 386 poachers have been arrested during 2014 in South Africa, and 317 poachers were arrested during 2015. Of these 317 poachers, 202 arrests were made in the KNP and 115 close to the park [27]. The area with the biggest problem (KNP) has a very small number of poaching arrests in relation to the number of poaching attacks. Clearly the current approaches are not sufficient.

Rangers have to patrol large areas of game reserve and cannot patrol the entire area effectively and efficiently. A costly solution would be to employ more rangers. An alternative solution is to reduce the area rangers need to patrol. Our hypothesis is that by providing a high-level prediction tool to rangers, rhino poaching can be reduced by predicting the area where a next rhino poaching attack is likely to happen. This study investigates the suitability of a causal network to aid in mitigating the rhino poaching problem. The approach is to divide the game reserve into manageable cells, and to assign a poaching event probability to each cell.

Different sources of information are fused to obtain a statistical heat map (conditional probability distribution) of the KNP that shows the area where a next poaching event will take place. The different elements that are fused are explained in detail in Chapter 6, and include variables such as the proximity to water sources and the moon illumination percentage. The data are encapsulated into conditional probability distributions (using probability density estimation techniques) conditioned on the context, which includes the time of day, moon illumination, and weather.

There are mainly three types of crime analyses to be performed: tactical, strategic, and administrative [28]. Tactical crime analysis examines the short-term and is more reactive than predictive. It asks the question, “What is taking place now?” and sets out to put an end to it. It normally involves one suspect with multiple targets, or one target with multiple suspects. This can be translated to this problem as one rhino having multiple threats, as well as being a current and ongoing problem in South Africa.

Strategic crime analysis is aimed more at long-term and ongoing problems. It is more concerned with pinpointing areas having high crime rates and discovering ways to decrease those crime rates. Administrative crime analysis is concerned with the allocation and positioning of resources (such as the police, army, and anti-poaching units).

The military has a slightly different view of these levels in the sense that they operate on a strategic, operational, and tactical level. Tasks are created and understanding of the factors shaping the enemy’s actions are studied at the strategic level [29]. Enterprises are employed at an operational level, and the completion of tasks occur at a tactical level. The clear difference between the levels exist so as to translate political aims into goals that can be attained on the operational and tactical levels [29]. The scope of this thesis can be seen as a combination of all these analyses and levels: we want to prevent poaching in both the short term and long term, and we want an optimal allocation of resources.

1.3 RESEARCH OBJECTIVES

1. To reduce the area game rangers have to patrol
   Rangers need to patrol large portions of land and cannot always be where they are needed. Employing more rangers is very costly. The alternative is to reduce the space that rangers need to patrol so that they can be effectively deployed at poaching hotspots.

2. To predict the area where poaching events are more probable (being proactive)
   Predicting the area where poaching events are more probable will aid rangers to intercept a poacher before another rhino is attacked. It is better to be proactive than reactive: being reactive may only result in the apprehension of a poacher, without saving the life of a rhino.

3. To determine the suitability of the approach in increasing the survivability of rhinos
   The tool can be tested in a real-world situation to test its practical applicability.

4. To facilitate a transdisciplinary and collaborative (co-created) solution
   In this project, a transdisciplinary approach is seen as an approach where the project leads realise that the project cannot be completed by a single disciplinary team. This is not always easy in the science and technology sector.

5. To shift the understanding of the rhino poaching problem
   The rhino poaching problem is highly complex and not very well understood. An objective of the study is to shift the problem to a space where researchers can start to address the inner workings of the problem.

1.4 THESIS STATEMENT

In light of the above, the thesis statement of this work is that by providing a high-level poaching hotspot prediction tool to rangers, rhino poaching in the KNP can be reduced by predicting the area where a next rhino poaching is likely to occur.

As most of the current approaches are reactive, the aim of this study is to investigate if a proactive approach will lead to a reduction in poaching and not only the apprehension of poachers after the crime. This implies that reducing the surface area that rangers need to patrol is preferred over employing more rangers. The manner in which this is achieved is by developing a co-created, structured causal model that contains all the available information regarding poaching events in the KNP. New information is added to the system and predictions can be made. A poaching event probability is assigned to every cell of the map and the end result is a heat map that can be used to allocate resources such as rangers or unmanned drones. The aim of the predictive tool is to assist park rangers to concentrate their patrols at the areas where future poaching events are most probable.

1.5 DELINEATION AND LIMITATIONS

In this study, only rhino poaching is considered. Other forms of wildlife crime, especially the poaching of other animals, are outside the scope of this study. Furthermore, only white rhinos and the habitat preferences and grazing preference of white rhinos in the KNP are considered. Black rhinos are extremely shy and hard to locate. They are also much more aggressive than white rhinos. Black rhinos are browsers where white rhinos are grazers, which means that the two species' foraging habits differ substantially, which will also have an influence on their habitat preference, amongst other things.

This study currently only considers activities taking place within the KNP. It is true that rhino poaching is part of a complex socio-economical problem, and together all these variables form a complete picture of the situation, but for this study the problem is confined to the KNP. As far as data are concerned, only successful poaching event data are used where historic data are mentioned and used.

1.6 DEFINITION OF TERMS

- “Poaching” refers to the action of illegally removing horns from rhinos, often resulting in the death of the rhinos.
- “Poachers” refers to the individuals who commit poaching attacks.
- “Rhino horn”, as already stated, refers to what the normal person would point out to be the rhino’s horn, even though it is not really a horn.
- “Dehorning” refers to the legal removing of rhino horns either by the owners or authorised game reserve personnel.
- “Expert” refers to an individual with above-average knowledge in a certain area and is used interchangeably with “stakeholder” in this thesis.
- “Rhino” refers to white rhinos only, unless explicitly stated otherwise.

Other terms and (statistical) concepts will be discussed in further chapters.

1.7 UNDERLYING ASSUMPTIONS

1. Assume that a measure of corruption or ill-discipline is present in the KNP, although the origins and realisations of this are not known to us. Corruption changes the system, as the guardians are then part of the problem. Ill-discipline has the same effect, although it might not be intentional or malevolent as is the case with corruption. This must be reflected in the effectiveness of the rangers doing their job.

2. Assume that the development of a tool for decision making will give the rangers an advantage over the poachers who do not have access to the insights of the experts. This impacts the confidentiality during the development of the tool.

3. Assume that rhinos avoid camps and busy areas and that the locations of the rhinos are unknown. According to the literature it can safely be assumed that rhinos prefer to steer clear of camps and busy areas, although there has been the odd sighting of a rhino being “where he is not supposed to be”. There are very few rhinos that are collared or tagged, thus it was decided to assume that none of them are collared or tagged. If the rhinos were collared or tagged there would be data available, which there are not. Rhinos are usually not collared because there are too many of them and the size and shape of their necks present challenges as specific collars have to be developed and purchased. Knowing where the rhinos are also does not necessarily solve the poaching problem: there are many more factors at play than just knowing where the rhinos are.

4. Assume that rhinos mainly choose their locations based on the proximity to water. Rhinos drink water roughly two to three times a day in the warmer months, and at least once a day in the colder months, thus areas close to water will have a greater chance for having a rhino present. Rhinos are not herd animals, but mothers and calves will graze together until the calf is old enough to be on his own. Having a small calf forces the mother to be within a few hours’ walking time of water, as well as shelter.

5. Assume that poachers avoid camps for fear of discovery. Discussions with experts have confirmed that poachers prefer to avoid busy areas. This should be reflected in the tool by taking note of the proximity to infrastructure that can act as static deterrents to poaching.

6. Assume that, for the purposes of this study, the researchers and decision makers in the KNP do not have control over external factors such as law enforcement in neighbouring countries, market factors (supply and demand), and other socio-economic factors. Although this is not strictly true, it falls outside the scope of this study and are addressed in other studies and efforts.

7. Assume that there exists a reasonable, patrollable grid cell size that is adequate for incorporating the necessary spatial variables. This cell size needs to be decided during a workshop and from literature reviews.

8. Assume that it is reasonable to capture the causalities between ecological factors contributing to poaching events in a model in order to bring a logical structure and shared narrative to the

problem space. Further assume that causalities can be represented with appropriate probability distributions derived from expert knowledge and literature reviews.

1.8 SIGNIFICANCE OF THE STUDY

In this study the problem under consideration is one of humans versus animals (poachers versus rhinos), where other predictive policing applications study humans versus humans (criminals versus victims). This study is also, to our knowledge, the first BN application to wildlife crime. In the literature there are many applications of BNs to environmental modelling, and there are a few mathematical applications to urban crime, but no mathematical or BN applications that attempt to mitigate wildlife crime.

The novelty of the work lies in the approach that was followed, as well as the distinct elements present. This is the first BN model of the rhino poaching problem, and by collaborating with experts a deeper understanding was created concerning the complexity of the problem. The highly complex rhino poaching problem was shifted from an intractable space, to a space where it can be addressed and, possibly, resolved.

The contribution of the study is to show that a high-level prediction tool for rangers can reduce the space that rangers need to patrol, and in the end, attempt to mitigate the rhino poaching problem. The outcome of this study could be used to provide decision support to park officials and rangers with the necessary infrastructure and provide a general template for resource poaching mitigation.

1.9 BRIEF CHAPTER OVERVIEWS

In Chapter 2, the literature is discussed that addresses what is known about this problem, what types of systems exist, and if the various problems have been solved. The approach of this study is put into context within the literature and the gaps addressed by the study are highlighted. This chapter highlights the study's value by discussing what has been done and what has not been done.

Chapter 3 describes the most important theory and concepts that are needed in the rest of the thesis such as the concept of probabilities, Bayesian Networks, and the difference between likelihoods and probabilities. In Chapter 4 the research methodology is described in detail. This forms the roadmap of the thesis in terms of what is going to be done and how it is going to be done.

Chapter 5 and Chapter 6 form the core of the thesis. Chapter 5 describes the expert workshop as well as the model that was evaluated and validated by the experts. It shows the evolution of the model from a “prior perspective” BN to a comprehensive and realistic model. Chapter 6 describes the latest model in detail from the structure to the nodes, to the causal links and the probabilities.

Chapter 7 shows how to use an open-source Matlab toolbox for developing a BN. A Microsoft Excel template is also presented that is used to develop a BN from scratch. In Chapter 8 the evaluation and validation of the model is discussed and what-if analyses are performed to check the model’s alignment with the experts’ knowledge. Chapter 9 concludes the thesis with a discussion of the results and future work.

CHAPTER 2 LITERATURE REVIEW

2.1 INTRODUCTION

This chapter positions the study within the literature, and highlights gaps in the literature that are covered by the study. The chapter starts off with the theory base in Section 2.2 which describes the use of Bayesian Networks and expert elicitation. Section 2.3 details the applications of Bayesian Networks to environmental modelling and urban crime. Bayesian Network applications to environmental crime applications currently do not exist and this is a gap addressed by this study. Section 2.4 focusses on rhino poaching and discusses the differences between this study and the only other two rhino poaching studies that were published in 2015. Section 2.5 concludes the chapter by identifying the gaps in the current literature that is addressed by this study.

2.2 THEORY BASE

In this study the aim is to develop a predictive model for rhino poaching using a Bayesian Network. Owing to an initial scarcity of credible data, an expert knowledge approach is followed. The theory base is the theory behind Bayesian Networks as well as the elicitation of expert knowledge.

2.2.1 Bayesian networks

Bayesian Networks (BNs) are graphical models that consist of nodes (variables) and edges (arcs) between the nodes. The edges have direction, thereby indicating causal logical links between nodes showing relationship. For example, if there is a link from A to B, then A is a parent of B and B is a child of A. BNs are directed acyclic graphs (DAGs) because the edges have direction and there are no feedback loops in a network. Figure 2.1 presents an example of a simple BN. There are three nodes in this network connected by two edges. The network is directed (the edges have direction), indicated by arrows and there are no feedback loops, thus the network is acyclic.

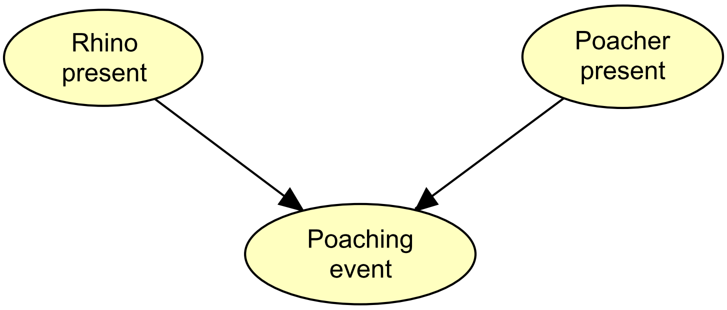

Figure 2.1. Simple rhino example

Any combination of expert opinion, empirical evidence, or literature is allowed in BNs [30–34]. Data are usually scarce, missing, and incomplete, nowhere more so than in the rhino poaching problem. BNs can be developed and populated with expert knowledge [34–37], but also with a combination of expert knowledge and empirical data in a single model [38, 39].

One of the most significant aspects of Bayesian methods is the concept of prior knowledge [38]. Priors demonstrate what is known about a problem before any research or experiments are undertaken. Prior information is clearly identified, decisions are made, and results are observed. These results are not to be seen as the final outcomes, but rather as new priors for the next round of experiments [40]. After the research is conducted, the priors are updated, and the updated version becomes the prior in a new iteration of the problem [38, 39].

Developing models forces researchers to think clearly and objectively about the problem at hand, which in itself is very valuable. Uusitalo [38] perfectly summarises the concept of BNs when she says, “Bayesian networks use probability as a measure of uncertainty: Beliefs about values of variables are expressed as probability distributions, and the higher the uncertainty, the wider is the probability distribution. As information accumulates, knowledge of the true value of the variable usually increases...”

The purpose of a BN is to encapsulate relationships between different variables and to interpret how changing a state in one node causally influences the other variables [41]. BNs are a way to capture and summarise information about an environment as well as the underlying assumptions behind a problem [30, 33]. BNs consist of a qualitative and a quantitative part: the qualitative part consists of the model structure, namely the nodes and edges. The model structure can also be determined quantitatively through structure learning algorithms. The quantitative part consists of the strength weights of the causal links between nodes [41]. A BN is a simplified representation of reality and is a means to facilitate active participation in solving the problem. The graphical manner in which BNs are represented is especially beneficial to experts or non-technical individuals who are not proficient in statistics. BNs act as a way to focus the attention and inputs of experts in an expert elicitation context.

Once a BN is completely specified it can be used for inference. Bayesian inference is where outcomes are analysed according to which node is queried and which nodes are observed [41, 42]. There is a variable's state that we wish to determine, but we cannot determine this state with certainty, thus we settle for computing how likely the variable is to be in a certain state [43].

McCann [39] states that BNs are not meant to be perfect representations of reality. The outcomes and predictions will be imperfect as is the case with any model, but the authors maintain that these models are still powerful in that they can represent complexity, causality, uncertainty, and variability in an understandable and intuitive way. Interactions between variables are presented as probabilities, and that makes BNs' risk estimation and uncertainty estimation better than those from methods that only take into account expected values [38].

The way ecologists express their results (such as the likelihood of their data given their hypothesis) is however of little value in communicating with stakeholders and managers [40]. Non-technical individuals, and especially those with non-statistical knowledge do not know or understand what it means to reject the null hypothesis at a p-value of 0.05. Researchers generally find it difficult to describe their certainty in their results or conclusions. Bayesian inference requires prior probabilities to be explicitly assigned, based on existing information, to the results of the experiments. These results can then be used to calculate posterior probabilities of the hypothesis given the data.

Expert knowledge elicitation unfortunately comes with its own challenges, such as getting the experts together, facilitating discussions, overcoming cognitive biases, and, the most challenging of all, obtaining useful probabilities and information. Expert knowledge can be used as an alternative source of data that can later be compared to the empirical data when new estimates become available. Expert knowledge can also be used as prior knowledge in the model and can be updated with new information [36, 44, 45]. The next section describes the work that has been done on the problem of knowledge elicitation.

2.2.2 Expert elicitation

Experts hold a wealth of contextual information that can be very useful in situations where little or no data are available. Applications of expert elicitation are varied and range from species richness estimates for coral reefs [44] and water quality modelling [46], to predicting the presence of the brush-tailed rock-wallaby [47], amongst others.

Elicitation can happen face-to-face or through interviews, phone calls, *etcetera*. Face-to-face interviews or workshops are more advantageous as it eliminates communication errors. This method also makes it easier to recognise when an expert does not understand something so that certain aspects can be explained to them in a different way [44]. Experts might be highly rated in their respective fields, but that might not translate into providing accurate and reliable probability assessments [44]. The value of the elicited knowledge depends on the expert's level of expertise, and ability to be objective.

Eliciting knowledge, especially probabilities, from experts is a challenging task. Uusitalo [38] states that one of the challenges of working with BNs, especially with expert knowledge, is to extract knowledge from experts in such a way that it can be turned into probability distributions. Experts find it difficult to provide probabilities without data, and some feel uncomfortable thinking about distributions instead of point estimates and confidence intervals [35, 38]. It could be that they fear being misrepresented (such as in an academic publication or a news article), or that they are unsure of how to quantify their own knowledge. Instead of supplying hard numbers, experts prefer to adjust a certain empirical estimate up or down [44], which is what was done in this study. This elicitation technique will be described in Chapter 5.

The following five papers illustrate the difference between eliciting knowledge directly and eliciting knowledge indirectly. The direct or structural approach to expert elicitation is the easiest approach from the researcher's perspective [45]. This entails asking the expert directly for values corresponding to their belief in parameters. An indirect approach is followed when the experts are not familiar with probabilities. James *et al.* [45] believe that an indirect approach allows for greater accuracy. Examples of where the indirect approach is followed include the papers by James [45] and Fisher [44].

Both James *et al.* [45] and Fisher *et al.* [44] develop software tools for eliciting expert knowledge, named *Elicitator* and *ElicitN* respectively. The tool developed by James *et al.* [45] shapes the quantification of expert knowledge so that it can be used as a prior in Bayesian regression. The tool provides a graphical interface for the experts to share their knowledge, and provides quick graphical feedback, which is said to result in better elicitations. The expert is asked for the best estimate of the probability of a certain scenario [45]. The expert is also asked to provide upper and lower quantiles as well as estimates of the upper and lower bounds for the probability of that scenario. Lastly the expert has to supply a confidence rating for every probability.

Fisher *et al.* [44] use their software tool, *ElicitN*, to estimate species richness in coral reefs. The software tool helps to elicit probability distributions from the experts by capturing the expert's "best guess" as well as the variability present to calculate likely intervals. The authors state that the first step in designing an elicitation process is to develop a statistical model that represents the underlying model [44]. As with James *et al.* [45], the authors of Fisher *et al.* [44] elicit subjective probability distributions from the experts that capture their uncertainty. Experts were also asked for their "best guess" of the "most likely" value. This value was also elicited as the mode, and not the mean or the median. The experts could choose whether they wanted to answer the questions as a percentage or as a multiplicative factor. A total of six parameters are elicited from the experts.

Another example of where an indirect elicitation method was used is illustrated by Denham [48]. Their approach is called an "indirect predictive P-method of elicitation" and was designed specifically for ecological modelling. The method uses an interactive map-based Geographical Information System (GIS) tool. The tool is used to elicit site-based parameters for a regression model. The expert selects a subset of the available features in the software tool, as well as a specific point on the map. A pop-up dialogue box then presents the expert with the different variable states at that location (much the same as with the grid for the rhino poaching problem). The expert is asked to consider 100 similar sites and to then supply a value for the median. These values present a conditional probability table (CPT) for the nodes. The expert also has to supply a lower bound and an upper bound for assigned probabilities.

Van Houtven *et al.* [46] developed a protocol to combine ecological and economic processes, thereby integrating environmental modelling, expert elicitation, and "nonmarket valuation methods". They illustrate the protocol with a case study of nutrient loads to lakes and lake water quality in North America. Das [49] devises a strategy to populate the CPTs as well as lighten the knowledge acquisition load by using weights to quantify the strength of the causal links between parent nodes and child nodes. These weights are elicited from experts and are used in a weighted sum algorithm to populate the CPTs. The weighted sum algorithm lightens the load on the expert in the sense that the expert has to come up with fewer probability distributions than usual. This increases the consistency as well as the expert's confidence. This approach is very similar to what was done in this study.

In the next few paragraphs, papers are cited in which the process is described whereby CPTs are populated. Wang [50] developed a conditional probability tree (CPTree) as an attempt to simplify the elicitation process. It presents a node's conditional probability table (CPT) as a tree view where the expert can expand or collapse the list of nodes and states according to his preference. The author also developed a shrinkable conditional probability table (sCPT) that gives a different view of the CPTs. Both tools contain graphical tools for elicitation based on bar charts and pie charts. Developing a network with the help of experts is a daunting task, which is one of the reasons why a rhino poaching network structure was developed before presenting it to a panel of experts. Experts state their beliefs and encode those into probabilities. Sometimes, however, the probabilities can be data-driven and objective [50].

Pellikka *et al.* [51] use expert knowledge and infer causal relationships between variables to evaluate the effects of game management strategies on population figures of Finnish wildlife species. Interviews were held with experts and they then each built a separate BN. The experts defined variables and expressed their knowledge in terms of a probability distribution. This distribution was defined by what the experts believe the direction, and magnitude of change over the variables will be from the present time until 2020.

Albert *et al.* [52] combine the knowledge elicited from various experts in order to construct a subjective prior for a Bayesian model. The knowledge is updated with new knowledge that becomes available so that the posterior can form the starting point for a later analysis. The authors work with the combined probability distribution of the possible parameter values, and not with the probabilities of events. Two test cases are discussed, the first a food risk assessment and subsequent dose-response for contamination by mice, and the second a study on the time it takes PhD students to complete their studies.

Researchers who develop expert knowledge-based BNs sometimes ask almost the impossible from experts: accurate and near-perfect input. Hope *et al.* [35] developed two applications, namely *CPTable* (probability entry) and *Verbal Elicitor* (probability elicitation) to mitigate this challenge. *CPTable* allows for nodes to be customised and contains sliders for binary elicitation, and tables can also be filled automatically. *Verbal Elicitor* allows the entry of values in plain English. Instead of supplying probabilities, experts can choose options such as “certain” or “almost impossible”. These terms are then translated into probabilities by the tool. *Verbal Elicitor* does not require the user to know anything about probabilities or BNs.

Hansson *et al.* [53] developed a BN as part of a study performed on a nuclear facility. Owing to a lack of data, the authors developed a method to include expert opinion into CPTs. The authors looked at eliciting a single probability as well as eliciting a full CPT. When eliciting a single probability, the authors found *Probability scale* the easiest option for the expert to use. Experts are generally unwilling to convey specific probabilities, so it was found that the *Likelihood method* was the best to use. The *Probability scale* is a vertical or horizontal scale where the expert can make a mark according to his beliefs, and the distance is then measured from zero to his mark. This is beneficial especially for experts who are not familiar with probabilities. The *Likelihood method* is based on the idea that the expert does not need to provide probabilities for every single CPT entry. The expert supplies the analyst with values for each state of each parent node as well as influence weights. In short, the expert supplies the analyst with a baseline as well as weights for how to move this baseline up or down over the different states.

The sizes of CPTs combinatorically explode with an increase in the number of parent nodes, and especially in the number of parent node states. A number of papers have been written to address this problem. The problem of eliciting probabilities from experts in large and impractical CPTs is addressed in the paper by Kokkonen *et al.* [54] where the authors’ motivation is to more easily elicit probabilities from experts to populate CPTs in discrete BNs.

Kokkonen *et al.* [54] develop a method that describes the CPT in terms of link strengths of edges between each parent node and each child node. The link strengths range from -1 to 1, where 1 represents a perfect correspondence between a parent node and a child node. A link strength of 0 represents independence between parent and child, and a negative value indicates an inverse relationship between parent and child. The goal of the work by Kokkonen *et al.* [54] is to present a technique to explain the CPT with a number of link strength parameters. These link strengths can be elicited from experts.

Numerous research papers that address the problem of large CPTs use a form of NOISY-OR or NOISY-MAX [55]. These methods use a small subset of questions that are elicited from experts to populate the CPTs [30]. Spaccasassi [55] created a web-tool that combines findings by asking text-based questions that correspond to each entry of the CPT. Answers are then grouped together by scenario. Providing accurate numerical values as answers to questions is another difficult task for experts [55], but there exist methods to simplify this process. Richards *et al.* [56] have first-hand knowledge of how difficult it can be for experts to convert their knowledge and beliefs into probabilities for a CPT, especially if more than three parent nodes are associated with a child node. Before they developed the *App2Adapt* tablet application, the authors printed out CPTs and gave it to experts to write down their beliefs. The struggle that the expert experience is heightened by the fact that multiple CPTs need to be completed, thereby reducing the strength of the model.

2.3 BAYESIAN NETWORK APPLICATIONS

In Section 2.2.1 BNs were described as the theory base on which the thesis is built. This section details instances where BNs are used in different real-world applications. In 2007 BNs were not a common method in environmental modelling yet [38], but at that time it was increasing in popularity. Researchers started realising that BNs are well suited to model uncertain and complex fields such as in environmental research. However, BNs are not without flaws, and Uusitalo [38] summarises these points in her paper. She goes on to say that BNs have been applied to many diverse problems and subject areas, increasingly also in environmental problems (as will become evident in this section).

2.3.1 Environmental modelling

Pullar [57] modelled Koala occurrence in a coastal region of Australia. The coastal region is under threat from urban expansion, thus urbanisation and dog presence are two factors contributing to the

Koala's absence in certain areas. The authors used a BN to discover, analyse, and understand the relationships between the threats and Koala conservation.

Interviews with three experts were conducted to construct a preliminary network [57]. The experts were asked to populate the network with probabilities after the network structure was explained to them. After the elicitation, Bayesian learning was applied to the network so as to update it with the observed data of Koala sightings. The outcome probability was calculated for each cell in the grid, just as is done in this study with the rhino poaching problem. It was found that both habitat and disturbance have a causal influence on Koala presence. Both are latent variables, as they are not observed directly. Latent variables are also present in this study and will be discussed in Chapter 6.

Pullar [57] gathered habitat quality information by using vegetation and terrain data, and disturbance was deduced from urban density and dog presence data. They also established that a grid of 250 metres was the smallest area a Koala's home range could be, and that was the size at which the observed data were fused. The home range analysis grid was large enough to detect patterns. In this rhino poaching study a grid of 5×5 kilometre was used as that was a suitable size to capture the necessary spatial variables at a certain location. Home ranges of white rhinos could be studied and used to obtain a better estimate for grid size. Home ranges were, however, not used in this study as each group of rhinos have their own home range and there are thousands of rhinos in the KNP.

The aim of the project by Stelzenmüller *et al.* [58] was to construct a Bayesian Belief network (BBN) - GIS framework to visualise the connections between collective human pressures, delicate marine terrains, and terrain vulnerability. According to the authors very few studies have integrated BNs and GISs fully. The framework also evaluates the impact of likely marine planning goals, and maps uncertainty-related variations in management measures. According to the authors, the BBN-GIS framework is a viable tool permitting for (1) the visualisation of connections, (2) the spatial evaluation of uncertainty related to spatial management scenarios, (3) the engagement of various stakeholder opinions, and (4) the enabling of quick updates of new data and connections [58].

Stelzenmüller [58] needed a tool that could both visualise the complex connections, as well as evaluate scenarios under certain planning objectives. A marine spatial plan details the spatial and temporal assignment of resources, but it should also evaluate the accompanying uncertainty of the data. The plan should be able to visualise uncertainties associated with the results of possible scenarios. The authors chose to define their prior probabilities with GIS data, and train their BBN with geospatial data only. What-if analyses can be used to evaluate the model for different planning objectives, as was done for the rhino poaching problem.

Dlamini [31] developed a BN to establish which features impact wildfires in Swaziland by using wildfire data for the time interval 2001 to 2007. The BN is constructed from evidence from satellite, wildfire data, and geospatial data. Causal links between variables were established from the literature, domain knowledge, and from the opinions of local fire management personnel. The same path was initially followed for the rhino poaching BN where causal links between variables were established from literature and domain knowledge. According to Dlamini [31], predicting and analysing wildlife ignition calls for a binary dependent variable if the prediction is based on a Bayesian or similar modelling technique. In this case that dependent variable, fire, indicates whether fire occurred at a point or not. The author fused the data in the GIS software ArcView and loaded that into Netica (commercial BN software) for BN calculations.

The two most important parts of wildlife management planning are (1) to determine which areas have a high possibility of burning, and (2) an explanation of burning patterns and also ascertaining the effects at landscape level [31]. The continuous variables defined by Dlamini [31] are road density, distance to human settlements, livestock density, human population density, elevation, slope, and climatic data. The discrete or categorical were soil type and land cover types. Sensitivity analyses enabled Dlamini [31] to single out variables that contribute the most to wildfire incidents, and she could thus rank the variables according to their impact on wildfire patterns. This author states that not everything is known about the interactions of human activities, climate factors, or rate of wildfire occurrences. The sensitivity of wildfire patterns to changes in these factors is also unknown. This proves that there exists a need to integrate research and modelling tools.

Johnson *et al.* [59] present a new heuristic method, called an iterative BN development cycle (IBNDC), to generate a tool for merging available data with expert knowledge. The IBNDC consists of a core process and an iterative process [59]. The core process is central to the network and is executed when the modelling of a network starts. Experts need to interact with each other and target nodes need to be defined, crucial factors need to be identified, and subnetworks grouped. The iterative process consists of four iterative phases, namely definition, quantification, validation, and evaluation of subnetworks. The iterative process uses the processes of BN modelling software, together with the inputs of experts, to break down a big task into smaller, more doable parts. By using the BN in a predictive mode the authors could evaluate the BN through inference. After all the evaluations were completed, an external panel of experts evaluated the BN.

The aim of the study was to model the success of cheetah relocations with a BN and thereby increasing the survivability of cheetahs [59]. BNs were constructed during a four-day workshop attended by cheetah experts and statisticians in South Africa. The focus of the workshop was the success of a relocation event and the elements that impact it with regards to the long-term viability of wild cheetah populations. Success is quantified with respect to the short-term survival of the relocated cheetahs, and the long-term survivability of the population.

The aim of the study by Farmani *et al.* [60] is to pinpoint a groundwater conservation policy against groundwater contamination management of a Copenhagen, Denmark water supply. The authors state that the method they propose can be used as a decision aid tool for process modelling for complicated problems. The method is applied to a case study and the method successfully singles out inconsistencies in the developed BN and then creates many management options that have a good trade-off between different goals.

The approach proposed in the study by Farmani *et al.* [60] imparts a framework to merge various factors that influence the complex problems [60]. The framework also allows inconsistencies to be uncovered in the decision analysis process, as well as creating management options that brings about an optimal trade-off between contradictory goals. According to the authors, the evolutionary algorithms can be used as an additional tool in a participatory unified evaluation process. The authors state that by using their integrated approach, the action space can be analysed more openly [60].

Hamilton *et al.* [61] developed a BN using expert knowledge and limited and scarce data to model the habitat suitability of a rare breed of giant freshwater crayfish. They also show how different evaluation techniques can be used in spite of scarce data. The authors used expert knowledge in the form of surveys to supplement the lack of empirical data. The goal of the surveys was to elicit probabilities that describe the strength of the relationship between the habitat suitability and three variables that were identified as important (elevation, upstream riparian (relating to wetlands adjacent to rivers) condition, and geomorphic condition). Each of the three variables has three states, yielding 27 possible scenarios. The experts had to give a probability for each of these 27 possible scenarios for high or low probability of suitable habitat for the crayfish. In the rhino poaching study we simplified this process by only asking the expert to specify certain values and then populate the rest of the table, thereby decreasing expert fatigue.

Hamilton *et al.* [61] then developed 18 alternative model versions based on this model structure, but each parameterised differently (some on expert knowledge, some on data, some on both expert knowledge and data). These models were evaluated by comparing their respective sensitivity analysis and performance accuracy. Hamilton *et al.* [61] used a function in Netica to conduct sensitivity analyses on the different models. The evidence entered into the nodes is varied to simulate the habitat suitability outcome after which the probability distribution is recorded. The main conclusion from the study is that the data models and the data-expert models performed better than the purely expert models. The elevation seemed not to have a big influence on the occurrence of the crayfish. However, the geomorphic condition had the biggest influence of all the variables.

2.3.2 Urban crime

The 2008 paper of Baumgartner [62] presents a novel BN model for predicting the behaviour characteristics of offenders committing new single-victim homicides. The BN can deduce the characteristics of an unknown offender from crime scene data. This aids the police in reducing the list of possible suspects for a specific homicide. The BN is developed from a database of completed homicide cases, and is learned from historical homicide data [62]. The authors also present performance metrics to calculate "...the sufficiency of an available database without knowledge of the underlying joint probability distributions." These performance metrics determine the size of the training data. Baumgartner *et al.* [62] presented two homicide cases to experts. The experts correctly predicted 53% of all offender variables, while the BN predicted 86% of all offender variables. Variables that caused disagreement between experts were also correctly predicted by the BN.

In the study by Zhang *et al.* [63], the authors develop a dynamic BN (DBN) to model the behaviour of opportunistic criminals in urban areas to assign an optimal patrol strategy to the police. According to the authors, criminals in urban areas are opportunistic and use knowledge of the police patrols to their advantage. Opportunistic criminals are "less strategic in planning attacks and more flexible in executing them" [63]. Zhang *et al.* [63] wanted to learn how criminals choose targets and what the likelihood was of them committing to that target. In the study it is assumed that all criminals behave the same, and that all the patrol officers are the same. The authors list their contributions as (1) learning their DBN model from criminal activity data, (2) presenting a compacted version of the DBN, and (3) an iterative learning and planning tool that updates with new data. The criminal behaviour as well as the interaction with the police on patrol are represented in the DBN. The compacted version of the DBN results, according to the authors, in superior learning accuracy and speed. The authors state that experimental validation with real-life crime data has confirmed their assumptions and the choice of model.

2.4 RHINO POACHING

Border safeguarding shares similarities with the rhino poaching problem, as many of the poachers are from neighbouring countries [12] and use the Kruger National Park's (KNP's) border to enter South Africa. Border safeguarding units are deployed in the KNP to capture poachers [24]. According to a SANParks official, most of the rhino poaching incidents in South Africa occur in the KNP, with 70% of those poaching incidents occurring along the KNP's 4,000 kilometre border with Mozambique [64]. The key reason for these high poaching figures in the KNP is that more than 90% of the world's white rhino population and 40% of the world's black rhino population are found in South Africa [15], and most of these rhinos are found in the KNP. Another important reason is that the KNP shares a border with Zimbabwe and Mozambique [24].

The following subsections present studies performed on the rhino poaching problem. The reader's attention is focussed on the fact that no tangible proactive solutions are executed or proposed. Section 2.4.2 presents the only two known rhino poaching models that were published in 2015, and the differences between them and this study are subsequently discussed.

2.4.1 Addressing the rhino poaching problem

The paper by Milner-Gulland [65] investigates the links between different socio-economic aspects of the rhino poaching problem, such as financial gain, and the punishment for poaching rhinos and other animals in Zambia during the late 1970s and early 1980s. They study the motivation behind illegal exploitation of wildlife by developing a model to optimise the poacher's harvest under "imperfect law enforcement”, and use that to investigate the effects that different sentencing structures will have on poaching. According to Milner-Gulland [65], there is a big difference in how local poachers and organised groups react to law enforcement. Local poachers are deterred by local law enforcement, but organised groups can only be deterred through high-level law enforcement operations.

Methods to reduce incentives instead of regulating hunting are 1) to reduce the demand in consumer countries, and 2) to increase the costs of poaching by increasing wages elsewhere [65]. The person in charge cannot do anything about the demand or the economic climate in his country, but he can increase the cost of poaching by investing in local projects such as projects where a portion of the revenue from safaris are given back to the people of the community. For instance, Sabi Sand game reserve recruited and trained men from nearby settlements to become game rangers in order to empower the local communities as well as decrease the chance of poaching syndicates recruiting them first.

The conclusion Milner-Gulland arrives at is that it is better to have a penalty varying according to the output of the poacher, rather than having a fixed penalty system. The probability of the poacher being captured is also a critical factor in his decision to poach [65].

Leggett [66] argues that effective counter-poaching strategies can only be achieved by effective law enforcement. This is to be achieved by providing site managers with information that can be used to determine how to distribute resources effectively to improve management and protection of wildlife. According to the author, staff morale decreases significantly when animals are poached. This low staff morale in turn coincides with higher poaching levels. Poaching levels are reduced slightly when the morale of the rangers is high.

Although Leggett [66] focused their study on elephant poaching, the same principals can be applied to rhino poaching. The authors state that most of the Namibian parks’ resources are allocated to managing poaching activities through armed patrols, and providing the park with logistic support and the necessary infrastructure to sustain these patrols. According to Leggett [66], there are two types of constraints when it comes to limiting the anti-poaching efforts: the one is related to political support, and the other to funding. Factors that limit the curbing of poaching include the lack of support from the government (both institutional and administrative), the lack of income for people from the local community, and the lack of logistical support.

Nepal experienced a complete turnaround in their conservancy strategy with increased anti-poaching efforts that led to a decrease in the poaching numbers (only one rhino was poached during 2003) [67]. They adopted a strategy called a “sweeping operation” where the patrol men would patrol thoroughly and then camp in hot spots where they knew rhino poaching is common. A further strategy change was to better educate the people from the local communities surrounding the park concerning the importance of rhino conservation and also the benefits it could bring them. A law change in Nepal allowed army staff to legally arrest poachers outside the park, instead of just inside the park.

Martin [67] states that many countries experience an increase in rhino poaching due to insurgencies. The countries focus their military units on borders or other vulnerable areas, and extract their patrols from national parks. Poachers see this as a golden opportunity to poach wildlife, especially elephants and rhinos. This has been prevalent mostly in Africa, and somewhat in Asia. Other examples include Mozambique and Zimbabwe, two countries which border on South Africa and the KNP. Nepal increased funding for intelligence, ensured better collaboration between parks, recruited army and NGO (Non-Governmental Organisation) personnel, implemented a new patrol system, upgraded telecommunications, acquired help from neighbouring communities, and implemented better leadership.

Lockwood [68] investigates which variables (spatial and management-related), and which combination thereof, best describes how the rhino poaching incidents are distributed in KwaZulu-Natal (KZN). She also explores attitudes towards poaching and rhino protection costs through the use of questionnaires. From the quantitative analysis she concluded that there was a need for additional research within state-owned parks to obtain a better understanding of poaching patterns. Her work could not identify the property related factors that influence the poaching patterns in KZN, however, it does call attention to spatial concepts that should be investigated in future research. A similar approach could be used for poaching in the KNP.

Kahler [69] use focus groups and interviews in Namibia to identify poaching drivers which range from the traditional sustenance (“cooking pot”) and economic (“pocket book”) motivations, to other less-traditional drivers of poaching. The authors found that the individuals interviewed had very high awareness of the rules and laws concerning poaching and wildlife crimes. This is in stark contrast with previous studies where researchers found that communities in developing nations had a low level of knowledge about wildlife laws. According to Kahler [69], the fear of getting caught as well as the effectiveness of the law enforcement both play a vital role in addressing poaching. Detecting and deterring requires that individuals get involved. Furthermore technical approaches such as surveillance and microchipping animals can also assist. Kahler [69] cite Rowcliffe [70] who concludes that wildlife crimes could be reduced by "increasing the probability of detecting violations rather than increasing penalties". This is in agreement with Milner-Gulland [65].

Kahler [69] repeatedly state the importance of transdisciplinary collaboration in managing the risks associated with wildlife crime such as poaching. This is especially vital in areas where wildlife populations are in close proximity to humans (such as rhinos in close proximity to impoverished communities), and penetrable borders of regions that cannot put a stop to illegal trading. Kahler [69] is in agreement with Milner-Gulland [65] that a static one-size-fits-all approach to punishment will not yield satisfactory results. In this thesis we also approach the rhino poaching problem from a transdisciplinary perspective, as it became clear early on that the problem is multi-layered.

In his 2012 article Eloff [12] examines the extent of rhino poaching incidents in South Africa by using GIS and remote sensing. Multiple spatial analytical techniques are incorporated with research to address the problem in 12 steps. The research was performed in the KNP from January 2010 until May 2010 when 71 rhinos were killed. Eloff [12] concludes that the rhino poaching problem in South Africa is highly complex and consists of a web of activities of poor neighbouring communities as well as corrupt officials. He lists suggestions for possible ways to curb this problem, some of which have already been implemented, such as imposing strict prison sentences for offenders in order to deter future poaching events.

Duffy [2] conducted a study of the literature and found that the relationships between poverty, poaching, and trafficking is not well-researched. They state that poverty is both directly and indirectly linked to poaching and trafficking of wildlife contraband, especially ivory and rhino horn. There are also different types of poachers which have their own way of poaching. Different types of poaching will require different policy responses. Poaching and trafficking are not necessarily driven by poverty, but more likely by wealth.

The main policy responses as summarised by Duffy [2] are, (1) changing people's behaviour through incentives (both positive and negative), (2) development of tourism as a means to reduce poverty, and (3) legislation of ivory and rhino horn internationally. Negative incentives include campaigns to change public opinion about the problem (such as campaigns to save the rain forest), and also enforcement designed to stop poaching activities. During the 1980s in Tanzania, the budget for law enforcement, and counter poaching patrols, were increased, leading to a successful decrease in poaching. This led to the rhino, elephant, and buffalo populations recovering. The disadvantage of enforcement, however, is that it can undermine the relationship between law enforcement officers and the community [2]. The community members perceive the law enforcement officers as taking away a resource that is irreplaceable to their lives and their livelihood.

Positive incentives could include financial benefits for communities. This will involve integrating the local community into conservation efforts to make them part of the solution. Locals could become part of tourism initiatives, trophy hunting, and the sale of goods. Distractions include alternative options to help distract communities from poaching and make poaching seem less appealing. Duffy [2] state that to understand poaching, we must understand human decision making. The individual does not decide on his own to poach, as he is influenced by many socio-economic and political factors. The authors further state that the weight of different poaching factors differ according to region and levels of wealth.

According to Duffy [16], the “Big 4 range states” are Zimbabwe, Kenya, Namibia, and South Africa. Combined, these four countries are home to almost 99% of Africa’s wild white rhino population, and about 96% of Africa’s wild black rhino population. Duffy [16] performed a quick review on the main variables in rhino conservation and threats to rhinos in these four states. According to this report, the main threats to (both white and black) rhinos are poaching, disinvestment, and insufficient resources. Rhino poaching is driven by the illegal demand for rhino horn by overseas countries. Disinvestment occurs because it is becoming more expensive and risky to protect rhinos. The resources are currently inadequate to protect the rhino populations [16]. In South Africa, the biggest threat is poaching combined with possible disinvestment by private rhino owners. The incentive to protect rhinos is decreasing rapidly.

The main policy that was suggested for South Africa is that of increasing incentives for private rhino owners so that they can continue to invest in rhinos and that the rhino population can grow and flourish [16]. Another important policy is to address corruption and organised crime. Intelligence sharing should be increased and captured more efficiently to enhance the analysis of such intelligence information. Ongoing training and development should be provided to assist in more arrests and convictions, and better levels of investigators and investigations. Alternative approaches to anti-poaching should also be investigated (such as the predictive modelling approach discussed in this thesis). Cross-border collaboration with Mozambique can also assist in curbing poaching.

According to Duffy [16], the plan for South Africa's rhinos should not rely on reactive approaches (prosecutions after the rhino has been killed), but rather strengthen proactive approaches by developing better systems to prevent poaching attacks. This is where the predictive model described in this study can play a big role. Relationships between neighbouring countries are improving and leading to tip-offs to aid rangers in stopping a poaching attack before it happens. Duffy [16] also state that game reserves and agencies have started to investigate the use of new technologies, and the improvement of existing technologies. Drones are marketed to game reserves, but Duffy [16] believes that they are inadequate for the task.

2.4.2 The latest in anti-poaching modelling and decision support

As can be seen by the above literature review, few studies exist to compare with this study. Most studies either focus on the financial impact of rhino poaching, decreasing the incentives for poaching, or the purely ecological impact. Two recent papers that were published in 2015 by Park *et al.* [71] and Critchlow *et al.* [72] come close to addressing the same problems as in this project. The paper by Park *et al.* [71] presents an anti-poaching engine that is a special case of spatio-temporal optimisation problems whereby poacher (and rhino) behaviour models are given to the engine as input. The paper by Critchlow *et al.* [72] investigates the spatial patterns of various illegal activities concerning wildlife.

Table 2.1 compares the assumptions made in this thesis with the assumptions made by the authors of Park *et al.* [71] and Critchlow *et al.* [72]. If the assumptions made in the thesis are directly linked to a variable in the model, the name of that variable is shown in brackets after the relevant assumption along with the section where it is explained in Chapter 6. The entry is left blank where no corresponding assumptions are made in a study.

The papers by Park *et al.* [71] and Critchlow *et al.* [72] were chosen because these are the only known studies encapsulating the work done in the thesis to a certain degree. It is brought under the attention of the reader that this study's framework, as well as original model, was published in July 2014 by Koen *et al.* [25], while Park *et al.* [71] and Critchlow *et al.* [72] were only published during the second half of 2015.

This study is in agreement with Critchlow *et al.* [72] on the notion that collaboration and cohesion are required to fully understand the problem. The drivers and patterns behind poaching also need to be understood before any attempt can be made to mitigate the problem. As with this study, Critchlow *et al.* [72] initially thought that ecological factors (such as distance to water and distance to roads) are important for predicting poaching locations. From the work performed by Critchlow *et al.* [72] it became evident that ecological factors are not the most important factors. Knowing where previous poaching events occurred seemed to be much more important. The study of Critchlow *et al.* [72] focusses mainly on getting to grips with the drivers behind illegal activities. Ranger patrol routes are modelled to ensure effective law enforcement.

In contrast, the study of Park *et al.* [71] seems to overlap with this study by also predicting poaching locations. They do, however, make a number of questionable assumptions, which could hinder their progress in the future if they were to use their methodology and technology on different resource plundering research environments. They make use of drones, and assume that the presence of the rhinos being guarded is known. Drones are expensive and many drones will be needed to cover a park the size of the KNP, for instance.

The first similarity between this study and the work by Park *et al.* [71] is that the studies are both centred around protecting rhinos. The difference, however, is that this study focusses specifically on white rhinos. The reason for choosing to study white rhinos is that the behaviour, vegetation preference, and movement model of white rhinos are vastly different to that of black rhinos. The white rhino population is much larger than that of black rhinos, and white rhinos are less shy than black rhinos, thereby increasing sighting opportunities. It would not make sense to simply group the two subspecies together under the common heading of "rhinos".

Another similarity between the two studies is that both parties investigate ways to mitigate rhino poaching attacks through prediction and thus attempt to pre-emptively safeguard the rhino population. The maps of both the respective game reserves or parks are divided into grids with cells of equal length. For this study a grid size of 5×5 kilometre is used, whereas in Park *et al.* [71] a grid size of 400×400

metre is used. The reason for this substantial difference lies in the vast size difference between the two environmental areas. This study focusses on the KNP in the north of South Africa, spanning a very large area, roughly the size of Israel. The park used in the study of Park *et al.* [71] is Olifants West in South Africa, spanning an area of about 155 km^2, whereas the KNP is 19,485 km^2. The grid sizes of both studies are in line with the size of the corresponding game reserves or parks. The landscapes are also very different: there are roughly 35 different landscapes in the KNP, and there is simply not enough space in Olifants West to match the diversity of the KNP.

Ecological factors such as the distance to water, distance to buildings, distance to vegetation, the number of animals per cell, and the elevation is deemed as important by both this study and Park *et al.* [71] for the prediction of rhino poaching events. The study by Park *et al.* [71] also state that the distance to roads and the distance to freeways can play an important role in the prediction of poaching events. In this study it was initially believed that the distance to roads was both a positive and a negative for poachers. Poachers use the roads to enter and exit the park, but they also prefer to attack far from roads, as that is the best option for not being observed. Furthermore, it is also known that rhinos steer clear of roads. There are not freeways or highways close to the KNP, thus the distance to freeways is not a consideration in this study.

Other similarities between this study and Park *et al.* [71] are that poaching attacks usually happen during twilight and at night, and that poachers prefer to escape the park in darkness. Furthermore there is agreement that rhinos avoid humans and areas frequented by humans. The probability of a cell being under attack by a poacher is also calculated by both studies. Both studies use models which have spatial and temporal components, and Park *et al.* [71] divide their temporal aspect into hourly intervals. The study by Park *et al.* [71] uses a different spatio-temporal graph for each agent, whereas this study uses a different instantiation of the model for each cell.

The biggest difference between the two studies can be seen on the methodological level. The study by Park *et al.* [71] uses a special case of spatio-temporal optimisation problems together with behavioural models, as well as multi-variate regression, in contrast with this study's BN. The study by Park *et al.* [71] also relies on numerous assumptions and certain assumed sureties that decrease the likelihood to yield usable results in a different environment. 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.

Another big difference between the studies is that Park *et al.* [71] uses drones in conjunction with rangers for better patrols. This assumes that, if the study is to be adapted to other research environments (as they state it can easily be done), drones are permitted in those parks. The authors state that in the Olifants West study they had one drone for each ranger. A single drone can cover a relatively small area, but even for a park as small as Olifants West they needed six drones. Drones are also extremely expensive, and a high-tech military drone will cost around R1 million. If this study is to be adapted to the KNP, for example, there will be significant challenges. Drones are currently not permitted in the KNP. There are also a much larger number of rangers in the park than in Olifants West, thus the number of drones will have to be much larger. Such a large number of drones will impact the tourist industry of the park negatively, as tourists would not want drones flying above their cars, or scaring away the animals. The biggest downfall of this approach is the monetary implications. At R1 million per drone, covering the KNP will cost millions of Rands, and that is just the purchasing cost. The maintenance and operating costs will also be considerable.

This study as well as the study conducted by Critchlow *et al.* [72] understand that the patterns and the extent of illegal activities is key to solving the poaching problem. Both studies also agree that the drivers and the spatio-temporal variation of poaching is poorly understood, and that determining these drivers and patterns would bring forth better law enforcement.

The study by Critchlow *et al.* [72] focusses on identifying trend drivers behind different kinds of illegal activities, whereas this study only focusses on rhino poaching. Once again the research areas differ vastly: this study revolves around the KNP whereas the study of Critchlow *et al.* [72] focusses on the Queen Elizabeth Conservation Area in Uganda. The Queen Elizabeth Conservation Area is 1,978 km^2, which is roughly one tenth the surface area of the KNP. Both studies identify areas of greatest risk in terms of rhino poaching. Both studies divide the maps of the corresponding game reserves or parks into grid cells of equal size. As previously mentioned, this study divides the map of the KNP into 5×5 kilometres, whereas the authors of Critchlow *et al.* [72] divide the map of the Queen Elizabeth Conservation Area into 500×500 metres.

Both studies initially deemed ecological factors such as distance to rivers, terrain slope, wildlife density and land cover, to be important factors in predicting poaching events. Another point of agreement is the fact that areas close to water have a higher probability of having a high density of animals. Critchlow *et al.* [72] realised in their study that ecological factors, though important, are not the most important factors for determining poaching events. The location of previous poaching events is a better indicator of where future poaching events will occur. During this study the same conclusion was arrived at to a lesser degree. The main difference between the two studies occurs on a methodological level. The study of Critchlow *et al.* [72] explicitly models ranger patrol effort, whereas this study predicts the likelihood of poaching events.

Table 2.1. Comparison of assumptions

This study	**Park *et al.* [71]**	**Critchlow *et al.* [72]**
	Rhino assumptions	
Rhinos avoid camps and busy areas (*Proximity_to_static_deterrents*, Section 6.3)	Rhinos avoid certain cells	
Rhino locations are unknown (*Rhino_present*, Section 6.4)	Rhino locations are known	
Rhinos choose cells based on proximity to water (*Water*, Section 6.4)	Rhinos choose cells based on elevation and steepness	
Important to know the whereabouts of rhinos for poaching (*Rhino_present*, Section 6.4 and *Poaching_event*, Section 6.7)	Important to know where rhinos are for poaching	
	Animal behaviour and movement patterns stay the same and can be learned	
Rhinos have their own distinct behavioural patterns (*Historical_rhino_presence*, Section 6.4)	All animals share the same movement or behavioural patterns	All animals behave the same
	Animals behave according to a clock	
	Two animals move independently	
	Infer animal movement/behaviour from incomplete without	
Rhinos are the targets or victims (*Poaching_event*, Section 6.7)		

This study	**Park *et al.* [71]**	**Critchlow *et al.* [72]**
	Poacher assumptions	
Poachers do not know the working of the model		
	Poacher attacks one cell at a time only	
Poachers avoid camps (*Proximity_to_static_deterrents*, Section 6.3)		
Poacher locations are unknown (*Poacher_present*, Section 6.3)	Poachers follow a fixed, known strategy and prefer certain locations	
Poachers choose cells based on a variety of factors (*Poacher_present*, , Section 6.3)	Poachers choose cells based on water sources	
A poacher close to a rhino will try to kill it	A poacher will kill a rhino in the same cell if there is no ranger present	
Unable to model poacher behaviour completely	Can model poacher behaviour	
	Ranger assumptions	
Ranger corruption is present (*Corruption_index*, Section 6.5)	Zero corruption	Zero corruption
Rangers do not know where to look	Rangers do not know where to look	Rangers know where to go look
If not corrupt, a ranger will successfully stop a poaching attack if he is close to a rhino (*Active_deterrents*, Section 6.5)	Ranger will stop a poaching when he is in the rhino's cell	
	One ranger for every drone	
	Ranger and drone do not miss anything in a cell	
	Ranger can only monitor the cell that he is in	

This study	**Park *et al.* [71]**	**Critchlow *et al.* [72]**
	Ranger can defend animal in 10 minutes travelling distance	
		Rangers use GPS units correctly
Model assumptions		
5×5 kilometre grid is a good cell size	400×400 m grid is a good cell size	500×500 m grid is a good cell size
		Ranger patrol routes are calculated correctly
Slope is combined with the overall landscape ($Landscape\_preference$, Section 6.4)	Slope is an indicator of rhino presence	Slope an indicator of rhino presence
	Position of an animal at time $t+1$ only dependent on its position at time t	
The relationships between ecological factors can be modelled causally (the entire BN)		
Time of attack is independent of the cell where it occurs ($Time\_of\_day$, Section 6.3)	Time of attack is independent of the cell where it occurs	
Reasoned initially that distance to roads is both a negative and a positive for poachers, hence the variable was omitted from the model, but when the data was analysed it became evident that poachers preferred to commit crimes further from roads. ($Proximity\_to\_static\_deterrents$, Section 6.3)	Distance from a cell to the closest road is important for poaching events	Distance to roads influences illegal activities
Initially take ecological factors into account but	Ecological factors are important for	Ecological factors are initially thought

This study	**Park *et al.* [71]**	**Critchlow *et al.* [72]**
are unsure of the importance for poaching attacks	poaching	to be important
	Probability of poaching event occurring during early hours of darkness is Gaussian distributed	
	Presence of drones (always)	
	A drone can see any cell 1.6 kilometre from it	
		GPS units give correct readings
		Trends for 1999-2012 stay the same in future
	Drones can distinguish between humans and animals (rhinos especially)	
	Drones can distinguish between rangers, poachers, and tourists	
	Enough drones to cover all the rhinos in the park	
	"Attackable" cells are known	
Study can easily be translated to other smaller parks	Can easily translate this study to other larger parks	Model will work in data sets with frequent detections
	Drones are welcome/will be welcome in the park	
	Funding exists to purchase and maintain drones for an entire park	
	Google Maps/Earth can accurately calculate paths along any type of terrain	
	Agents can only move up, down, left, right (not diagonally)	

2.5 CONCLUSION

This chapter provides the reader with background as to why this study was undertaken. The theory base was discussed in terms of BNs and expert elicitation. Several BN applications pertaining to environmental modelling and urban crime were discussed before turning to the specific problem of rhino poaching.

The authors describing expert elicitation neglected to inform their readers how exactly they obtained the probabilities from the experts, or how they converted expert knowledge into probabilities. This process is explained in Chapter 6 for CPTs that are too large to elicit easily in its original form.

The literature review revealed that BNs have not been applied to environmental crime before, nor has a model been developed to address the rhino poaching problem. Rhino poaching has been discussed at length in terms of the economical and environmental applications in the literature. Policy changes have also been proposed to mitigate the problem, but none of them (up until 2015) seem to involve technical work such as mathematical models or predictive models. The conclusion is that there does not exist BN models for modelling or resolving the rhino poaching problem, which is the focus of this study.

A challenge that faces the KNP in terms of policing is the fact that the park is much larger than most other parks, specifically the ones mentioned in this chapter. There are also not enough rangers to patrol the area effectively. This in turn leads to not having enough high-quality data, and alternative methodologies have to be investigated. These methodologies are discussed in Chapter 4.

CHAPTER 3 THEORY AND CONCEPTS

3.1 INTRODUCTION

The human mind has preconceived ideas about how uncertainty works. There can be a number of different related events and outcomes, and the only way to make sense of these and to reach sensible conclusions, is to have a formal set of rules that enhances these ideas [42]. This set of rules consists of the concepts and the rules of probability that provide a framework of making sense of the world.

From Chapter 2 it follows that there are a number of concepts needed to address the rhino poaching problem. This chapter covers the important concepts and definitions that will be used throughout the thesis. These concepts serve as the building blocks for explaining Bayesian Networks. It will also aid in explaining the methodology used in this study.

3.2 PROBABILITIES

The simplest way to define a probability is as the proportion of times that a particular event occurs out of the total number of times that the event is repeated in an experiment [73], which is the frequentist approach to probabilities. An event could be anything from obtaining “heads” in a coin toss to throwing a six on a balanced die. An experiment is a repetition of events, such as 100 coin tosses or 500 die throws. When an experiment is repeatedly performed under precisely the same circumstances [74], the relative frequency or proportion of each event can be determined by counting the number of times that event occurred in the first n repetitions of the experiment ($n(E)$). The probability of the event [74] can then be written as

$$P(E) = \lim_{n\to\infty} \frac{n(E)}{n}, \tag{3.1}$$

which means that as the number of repetitions of the experiment approaches infinity, the frequency of the number of events E will be equal to the probability of that event [74].

The problem with this relative frequency treatment of probabilities is that it is unknown whether or not a precise and accurate number will be obtained for $P(E)$ given a finite number of experiments. It is also unknown whether or not the same results will be obtained if the experiment is repeatedly performed. Followers of the frequentist approach will argue that if n is infinitely large, the probability $P(E)$ will be precise and that Equation 3.1 should be used as an axiom rather than a theorem [74].

The biggest difference between the Bayesian approach to probabilities and the frequentist approach is that a probability in the Bayesian approach can be calculated for any event, while a probability in the frequentist approach can only be calculated on repeatable events. In the Bayesian approach one can thus, for instance, compute the probability that the moon is made of cheese and add evidence to the contrary, but no such calculation is possible in the frequentist approach.

If the events are *mutually exclusive* (meaning they cannot take place at the same time) and include all possible outcomes, then by definition the probabilities of these events sum to one. Probabilities are the cornerstones of any and all statistical methods. This definition of probabilities now paves the way for the more complex definitions of different types of probabilities, which will aid in understanding Bayesian Networks (BNs).

3.2.1 Joint and marginal probabilities

Bishop [73] starts off by supposing that A and B are two categorical random variables. The probability that A will then take on the value a_i and B will take on the value b_j can then be written as $P(A = a_i, B = b_j)$. This is referred to as the *joint probability* of A and B [73]. In the BN presented in this thesis only categorical random variables are used. Inherently continuous variables are discretised to simplify calculations. The problem, however, with discretising continuous variables is that there is a loss of granularity when partitioning them into bins.

Using the concept of joint probability, the sum rule is denoted by

$$P(A = a_i) = \sum_{j=1}^{N} P(A = a_i, B = b_j),$$

where $P(A = a_i)$ is the *marginal probability*. The marginal probability is obtained by marginalising (summing out) the other variable(s) (in this case B). In short this can be written as

$$P(A) = \sum_{B} P(A, B).$$

3.2.2 Conditional probabilities

Conditional probabilities are extremely important in the field of probability theory. Probabilities are often sought where there is partial information or evidence available concerning the outcome of the experiment. In that case the desired probabilities are conditional.

Assume that A and B are two events, and that B has already occurred. The conditional probability that A occurs given that B has occurred can be written as $P(A|B)$ and pronounced as "the probability of A given B". If B occurs, then for A to also occur, the occurrence must be a point in both A and B, thus $A \cap B$ ("A and B"). It is known that B has already occurred, thus B is the new reduced sample space [74]. The probability of $A \cap B$ occurring is the probability of $A \cap B$ in relation to B. The definition of a conditional probability as per Ross [74] can thus be written as

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{P(A, B)}{P(B)} \text{ if } P(B) > 0. \tag{3.2}$$

Multiplying both sides of Equation 3.2 by $P(B)$ yields

$$P(A, B) = P(A|B)P(B). \tag{3.3}$$

The definition of conditional probabilities leads to the derivation of Bayes' formula, which is a pivotal element of statistics in general, but more specific, of BNs. The derivation of Bayes' formula follows from Ross [74].

3.3 BAYES THEOREM

The product rule for A and B [74] can be written as

$$P(A = a_i, B = b_j) = P(A = a_i|B = b_j)P(B = b_j)$$

or

$$P(A,B) = P(A|B)P(B). \tag{3.4}$$

It can also be shown that:

$$P(B = b_j, A = a_i) = P(B = b_j|A = a_i)P(A = a_i)$$

or

$$P(B,A) = P(B|A)P(A). \tag{3.5}$$

Owing to symmetry, the property holds that $P(A,B) = P(B,A)$. Thus, equalling the right hand sides of Equation 3.4 and Equation 3.5 yields the following equation:

$$\begin{aligned} P(A|B)P(B) &= P(B|A)P(A) \\ P(B|A) &= \frac{P(A|B)P(B)}{P(A)}. \end{aligned} \tag{3.6}$$

Equation 3.6 is called *Bayes' rule*. "If we think of the events B_j as being possible "hypotheses" about some subject matter, then Bayes' formula may be interpreted as showing us how opinions about these hypotheses held before the experiment should be modified by the evidence of the experiment" [74]. Repeatedly applying Bayes' formula forms the basis of reasoning with certain or uncertain evidence [42].

3.3.1 The frequentist approach versus the Bayesian approach

The relative frequency approach that was taken at the start of this section to explain probabilities is a frequentist approach to statistics that assumes random and repeatable events, whereas the Bayesian approach provides a way of quantifying uncertainty. The Bayesian view is one that describes the uncertainty corresponding to the model parameters or even the choice of model itself [73]. "Bayesian subjective interpretation of probability sidesteps non-repeatability issues - it's just a framework for manipulating real values consistent with our belief about probability" [42].

Suppose we have an event such as a volcanic eruption. This is not a repeatable event that can be simulated in the lab multiple times to obtain a relative frequency. Certain information about the event will always be known, even if it is just educated guesses. For instance, it is known what the size of the volcano is. If new information is uncovered about, say, the density of lava or the speed at which the volcano erupts, the current knowledge about the behaviour of volcanic eruptions can be updated to incorporate this new information [73].

3.4 BAYESIAN NETWORKS

Bayesian Networks (BNs) are described in Chapter 2 as directed, acyclic graphical models that consist of nodes (variables) and edges (arcs) between the nodes. Nodes without parent nodes are called "root nodes", and nodes without children nodes are called "leaf nodes" [75]. Root nodes can be seen as the input variables of the network and leaf nodes are the output variables, although any node can be a target variable (output). Evidence can also be entered for any node [37]. Root nodes correspond to causes and leaf nodes to effects. Referring to Figure 3.1, there are three nodes in the network connected by three edges. The network is directed (the edges have direction), and there are no loops, thus the network is acyclic.

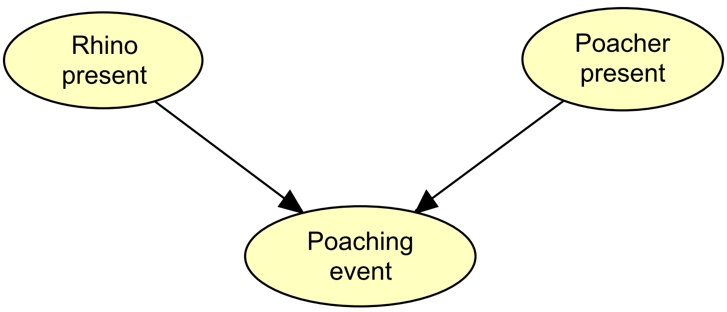

Figure 3.1. Simple rhino example revisited

A BN represents a "factorisation of a distribution into conditional probabilities of variables" that are conditioned upon parent variables [42]. Referring to Figure 3.1, the factorisation of *Poaching_event* can be written as $P(Poaching\_event|Rhino\_present, Poacher\_present)$. The joint distribution for the network can be written as $P(Rhino\_present, Poacher\_present, Poaching\_event) = P(Rhino\_present)P(Poacher\_present)P(Poaching\_event|Rhino\_present, Poacher\_present)$.

Each node has a table (specified by the modeller) that describes the prior probability of each state, or the conditional probability for each state if it is a child node. The posterior probabilities of states can be calculated using traditional Bayesian learning statistics. An advantage of BNs over other models is the graphical construction that visualises the relationship between variables more clearly.

The properties of the model, such as conditional independencies between variables, can simply be read off the graph. Intricate calculations can also be represented quite easily by using factorisation. The joint distribution of a node is conditioned on its parent nodes and can be written as

$$P(x_1, x_2, \ldots, x_D) = \prod_{i=1}^{D} P(x_i|pa(x_i)), \tag{3.7}$$

where $pa(x_i)$ are the parent nodes of x_i.

Variables in a BN can either be discrete or continuous, and discrete variables are represented by a finite number of discrete states [37]. The states must cover all possible scenarios and also be mutually exclusive. The states must also pass the “Clarity Test” [43] which states that someone who has all the information and data available about a problem, and can see into the future, should be able to understand the problem description without needing additional information. The BN can be called a discrete BN when all the nodes are discrete.

According to Uusitalo [38], the functionality of BNs lie in the fact that Bayes’ rule allows us to work in both directions in a BN. The probability distributions of child nodes given the values of their parent nodes can be calculated, as can the probability distributions of parent nodes given the values of their children nodes. Inference can be performed from causes to effects, or from effects to causes.

BNs are useful in facilitating stakeholder participation for planning and decision processes [39] and they can promote communication among ecologists, decision makers, and non-experts. BNs can depict and merge empirical data, graphically communicate the complex connections and problems, and address uncertainties in an organised way such as described in Chapter 2.

3.4.1 Conditional probability tables

The strength of causal links between nodes in a BN are defined as entries in Conditional Probability Tables (CPTs) [37], and the structure of the CPT comes from the factorisation of nodes. The CPT applies parent node states to those of the child node and includes entries for all possible combinations of parent and child node states. The CPT is represented as a table in the case where the parent nodes and child node are all discrete, as in the case of the simple rhino example.

CPTs rapidly increase in size and complexity with even just a small increase in the number of parent nodes [49]. If a child node has P parent nodes, and both the child and the parent have s states, the number of CPT entries will be s^{P+1}. For small numbers such as $s = 2$ and $P = 3$, the number of CPT entries already amount to 16 probabilities.

3.5 POSTERIOR, LIKELIHOODS, AND PRIORS

The conditional probability distribution of a variable can be written in words as

$$\textit{posterior} \propto \textit{likelihood} \times \textit{prior}. \tag{3.8}$$

The posterior probability is proportional to the product of the likelihood and the prior probability. The denominator in Equation 3.6 (Bayes' rule) is called the normalising constant. This guarantees that the posterior is a valid probability density function and that it integrates to one, hence Equation 3.8 is *proportional* and not *equal*. The prior probability captures the assumptions before observing the data [73]. The effect of the observed data is represented by the likelihood function which is a function of the parameter vector. It is important to note that it is a function and not a probability. The posterior probability evaluates the uncertainty in the parameter after the data are observed.

Writing Equation 3.8 as a formal mathematical equation yields

$$P(\theta|D) = \frac{P(D|\theta)P(\theta)}{P(D)},$$

where $P(\theta|D)$ is the posterior, θ is the unknown or hidden parameter vector, and D is the observed data [73]. The generative model or likelihood function is $P(D|\theta)$, and $P(\theta)$ is the prior belief concerning which variable values are appropriate [42]. The posterior is calculated after taking the observed data, D, into consideration. The most likely a-posteriori (MAP) setting [73] is the one maximising the posterior

$$\theta^* = \arg\max_\theta P(\theta|D).$$

For a so-called "flat prior" where $P(\theta)$ is constant and does not change with θ, the MAP solution simplifies to the maximum likelihood [42].

The frequentist approach also has a likelihood function, but it is defined differently to the Bayesian approach. In the frequentist approach the parameter vector is fixed and the distribution of all possible data sets D are considered [42]. In the Bayesian approach there is only a single data set and that data set was observed. The uncertainty is represented by a probability distribution over the parameter vector.

3.5.1 Probabilities versus likelihoods

Two terms that are commonly misinterpreted in statistics are "probabilities" and "likelihoods". In popular language, both refer to more or less the same thing: a number between zero and one attached to an event that encapsulates our belief in the occurrence of that event. How these two numbers are calculated, however, is where the biggest difference lies. Assume a very basic BN such as in Figure 3.2 consisting of just two nodes: *Poacher_present* and *Poacher_observation*, where *Poacher_present* feeds into *Poacher_observation*.

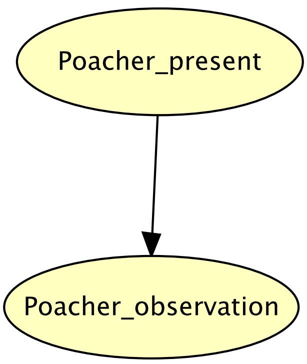

Figure 3.2. *Poacher_observation* example

From Equation 3.7 it follows that the conditional probability of *Poacher_observation* can be written as

$$P(P\_observation|P\_present) = \frac{P(P\_present|P\_observation)P(P\_observation)}{P(P\_present)} \tag{3.9}$$

where $P\_present$ and $P\_observation$ are used in the place of *Poacher_present* and *Poacher_observation* respectively for the sake of brevity.

Assume that each node has two states, namely "Yes" and "No". Further assume that the CPT is already populated and is presented in Table 3.1. *Poacher_present* is called a "hidden" variable as it cannot always be observed, and *Poacher_observation* is called an "observed" variable as it provides evidence.

Table 3.1. Example CPT for *Poacher_observation*

Poacher_ observation	*Poacher_present* No	*Poacher_present* Yes
No	0.9900	0.2000
Yes	0.0100	0.8000

The CPT can be read in two different directions, depending on which term is used, "likelihood" or "probability". Referring to Equation 3.8 and Equation 3.9, the quantity $P(P\_observation|P\_present)$ is referred to as the *posterior probability* and $P(P\_present|P\_observation)$ as the *likelihood*. Suppose we would like to know what the likelihood is that there will be a poacher present given that he is observed? The possible answers to this question are listed below.

$$\begin{aligned} P(P\_present = Yes|P\_observation = Yes) &= 0.8000 \\ P(P\_present = No|P\_observation = Yes) &= 0.0100. \end{aligned}$$

The values of $P\_present$ are varied while the values of $P\_observation$ are fixed because it is known that there definitely was a poacher observed. Note that the above equations do not sum to one, as this is a characteristic of likelihoods: likelihoods are functions and not probabilities. Likelihoods state how likely the observed data set is for different settings of the parameter vector.

Another question to ask is, what is the probability of observing the poacher given that he is present? The possible answers to this question are present below.

$$\begin{aligned} P(P\_observation = Yes|P\_present = Yes) &= 0.8000 \\ P(P\_observation = No|P\_present = Yes) &= 0.2000. \end{aligned}$$

The values of $P\_observation$ are varied while the values of $P\_present$ are fixed because it is known that a poacher is present. Note that the above equations do sum to one, which is in accordance with the laws of probabilities.

3.5.2 Conditioning versus dependence

Two variables are independent if, by knowing the state of one variable, does not provide any information about the other variable [42]. From a mathematical point of view: if the joint probability of two variables are equal to the product of the respective marginals, then the two variables are independent: $P(A,B) = P(A)P(B)$.

Suppose that A, B, and C are three variables such that $P(A|B,C) = P(A|B,C)P(B|C) = P(A|C)P(B|C)$. It can then be said that A is *conditionally independent* of B given C. This can be written as $A \perp\!\!\!\perp B/C$. For this to hold, A and B must be independent given all the states of C.

3.6 COUNTERFACTUALS

Counterfactuals are a way of reviewing unobserved potential outcomes of a situation [76]. It is counter to the facts and is concerned with situations or events that have already occurred. It can include a measure of hindsight such as “If only I had worn sunscreen, I would not have burned” or “If only I had my brakes fixed, I would not have rear-ended that car”.

Theorem 7.1.7 taken from Pearl [77] states that if the probabilistic causal model is M with a probability function $P(d)$ (where $P(d)$ is defined over the domain D), then the conditional probability $P(B_A|e)$ of a counterfactual sentence “If it were A then B” given evidence e, can be evaluated using three steps: (1) abduction, (2) action, and (3) prediction. In the abduction step the probability function

$P(d)$ is updated to result in $P(d|e)$. The action step modifies the model M by the logical action $do(A)$ where A is the antecedent (in this case it existed before the introduction of B) of the counterfactual. This results in the submodel M_A. The prediction step uses this submodel together with the new probability function $P(d|e)$ to calculate the probability of B which is then called the "consequence of the counterfactual" [77].

Counterfactuals can also be used when evaluating BNs. The more popular "what-if" analyses fall under this heading and its application to this model is shown in the results chapter later in the thesis. In counterfactual reasoning the leaf nodes of the BN are observed and the causes are inferred.

3.7 CONCLUSION

This chapter presented, and briefly described, the statistical concepts needed for presenting and describing BNs in later chapters. Probability theory underpins most, if not all, that we do, thus it is crucial to present and explain the concepts early on. The next chapter describes the methodology used in this study.

CHAPTER 4 DESIGN OF THE STUDY

4.1 INTRODUCTION

The previous chapter described important statistical concepts that are needed to understand Bayesian Networks, whereas this chapter introduces and describes the research methodology used for this study. The aim of the chapter is to provide the reader with insight of the thought process behind this study and to bring together everything as a whole.

Section 4.2 gives the overall approach of the work and presents the research methods that were used to design the study, and Section 4.3 explains important concepts used throughout the thesis. Section 4.4 describes the research methodology in detail. The research instruments and the data and expert knowledge are described in Section 4.4. Section 4.5 presents the limitations of the work, while the ethical considerations are discussed in Section 4.6 and the chapter is concluded with Section 4.7.

4.2 OVERALL APPROACH

The overall approach to this study was to develop a mathematical model to predict possible poaching event areas in order to reduce the patrol space. A mixture of research designs were used in this study; in this section each of them will be considered and their role in the study will be discussed briefly.

4.2.1 Survey-based research

At the start of the study, personal interviews were conducted with experts to gain knowledge about the rhino poaching problem. These personal interviews were informal and took the shape of a storytelling session rather than a structured interview. An expert knowledge workshop and further personal interviews were held to elicit expert knowledge. The aim was to evaluate and validate the first constructed model, but also to populate the model with conditional probabilities concerning certain occurrences and behaviour.

The reason for choosing interviews and a workshop was due to the lack of data and available knowledge in the field: the experts have a wealth of knowledge but this knowledge rarely translates to literature. The strength of this research design lies in extracting relevant information from experts by talking to them and letting them share their knowledge. In the paper by Johnson *et al.* [59] an expert workshop was held to construct and populate a causal model for cheetah relocation. The study would not have had the measure of success it had if it had not been for the inputs from experts.

The weakness of this research approach is the fact that human error is always present. Two different people might remember the same event differently, or be subconsciously influenced by each other to give the same answer [78]. It is also difficult converting the experts' knowledge into conditional probabilities since many experts, especially those who work in the field, are not necessarily adept at converting their knowledge into quantifiable variables.

4.2.2 Secondary data analysis

Secondary data analysis is the analysis of data that were previously gathered by other researchers. The reason that primary data were not used (or collected), is because the type of data in question is that of rhino poaching incidents. These incidents do not occur with regular time intervals such as the measurement of rainfall or temperature. Rhino poaching incidents are rare events (unfortunately not as rare as one would hope) and are only captured once a rhino is poached. A challenge in the rhino poaching problem, and most ecological problems, is the lack of complete sets of high-quality data. At the start of the study, rhino poaching data were scarce and thus an expert-driven approach was used. As the study progressed, the poaching data increased. The secondary poaching data were used to understand the problem and to ensure that the model is realistic. If the data were complete enough for the model, it would have been ideal to use a mixture of data and expert knowledge.

Throughout this thesis both data and expert knowledge will be mentioned. Geographical Information

Systems (GIS) data and literature sources were used for computing prior probabilities for different states of spatial variables, while expert knowledge was used to establish the conditional probabilities due to a lack of empirical data. Rhino poaching records were used to analyse the importance of certain variables and is discussed in Section 6.6.

The strength of secondary data analysis is that the researcher does not have to collect it. Assuming the previous researcher collected the data with care and precision it is much less time-consuming. If there are different datasets from different researchers to choose from, the current researcher can seek out the most experienced researcher's dataset. The paper by Warchol [17] used secondary data in an environmental problem to establish a picture of what the illegal wildlife market looks like in southern Africa.

The weakness of secondary data analysis is that if the researcher was not meticulous in collecting the data, the entire dataset could be obsolete. The data could also have been collected from a different perspective to the one the current researcher needs. For instance, the rhino poaching data were collected with the premise that the rhinos are the victims, when in fact they are the commodities. This argument will be explained later in detail.

4.2.3 Statistical modelling

The reason for modelling the rhino poaching problem is to understand and reason about the problem and to make predictions about it. If we understand the problem, we can reason and make predictions about it. Predictions are necessary in order to provide the operators at the KNP with rich patterns and a way in which to strike pre-emptively. If the structure of the rhino poaching problem was understood, different methodologies such as ontologies could have been used. Ontologies yield relationships between variables, but not causalities, whereas BNs offer the opportunity to use causalities to connect variables and to build understanding. However, ontologies may be useful for establishing the language and concepts that are relevant to the problem.

Modelling the rhino poaching problem by developing a Bayesian Network (BN) has potentially large theoretical significance by explaining the rhino poaching problem in terms of the drivers behind the problem. The output of the BN is probabilities of a poaching event given the state of certain variables. This also has practical implications by curbing rhino poaching.

The reason for using a BN is due to its ability to incorporate a mixture of data and expert knowledge, or even use only expert knowledge. BNs can learn and adapt as new data is integrated. Furthermore BNs are visually pleasing and therefore easy to explain to non-technical persons.

Statistical modelling's strength in this study originates from the strong statistical background on which the causal model is based as BNs encode probabilistic relationships among variables. The authors of Pullar [57], Johnson [59], and Borsuk [79] are but a few who successfully developed BNs to solve environmental and ecological problems. Borsuk *et al.* [79] use a BN to predict the effects of river rehabilitation undertakings and Johnson *et al.* [59] develop a BN to model the survivability of cheetah relocation in the wild. In the paper by Pullar [57], the authors construct a BN to model and represent the threats that koala populations experience close to urban areas. A weakness of this research design is the fact that there is no manner in which to establish how accurately the model reflects reality (except in hindsight). Complex relationships between variables can also be over-simplified or conversely lead to a computationally expensive BN.

4.2.4 Interdisciplinary research

The rhino poaching problem has many different facets, each situated in a different discipline. Mathematical models and statistical methods are merged with ecology and sustainability and researchers from a mathematical and engineering background are collaborating with ecologists, wildlife specialists, and law enforcement agencies. Collaboration between different fields and disciplines is established to create a more stable and viable space for resolving the rhino poaching problem.

Interdisciplinary research was used as it was found early on in the study that rhino poaching could not be viewed from only one area of expertise. Collaboration between disciplines is very important.

The strength of interdisciplinary research in the rhino poaching problem lies in the fact that multiple views on the same problem make a stronger case for a solution. The different disciplines and fields create a solid basis for the rhino poaching problem to be analysed. In the paper by Gonçalves [22] the case is made for interdisciplinary, multidisciplinary, as well as transdisciplinary research as a means to study problems such as the rhino poaching problem. Interdisciplinary research pertains to the transfer of knowledge to other disciplines but can lead to the forming of new disciplines. For instance, a software developer might look to a mathematician for help on a certain part of his problem. Multidisciplinary research pertains to the application of several different disciplines at the same time, while preserving each individual core discipline [22]. For instance, an ecologist working with a civil engineer to perform an environmental impact study before a new bridge is erected is seen as a multidisciplinary approach. Transdisciplinary research, as Gonçalves [22] puts it, "draws on existing disciplines and seeks to generate knowledge between disciplines to restore the 'weave' and thus move beyond disciplines." A challenge of this approach is that a skilled facilitator is needed to enable effective collaboration.

4.3 ADDITIONAL CONCEPTS

In the previous chapter most of the statistical concepts were explained. There remain a few concepts that are not of a statistical nature, but should be included for coherence.

4.3.1 Workshops and expert knowledge

An *expert* is an individual with expertise concerning the topic of interest [34]. This expertise can either be from research work or field work, but either way it is from personal experience. The term *expert* therefore includes individuals working at universities, or in the field. A *stakeholder* is an individual who either has the power to influence the topic or situation, or is influenced by it. Farmers might not have the power to influence environmental problems, but they are affected by it, and they are the people who will be implementing the decisions.

It is important to have both types of stakeholders present at a workshop as they help to clarify the problem space and provide differing opinions and views. They also increase the likelihood of the decisions being implemented after the workshop. Experts and stakeholders do not necessarily exclude each other, as an expert can be a stakeholder and vice versa. In some cases the analyst and the expert might be the same person [49]. The analyst develops the network using his own expertise, and then uses the network to make decisions. The terms *expert* and *stakeholder* will be used interchangeably in this thesis.

Experts are also rarely willing to give probabilities directly [35]. Eliciting probabilities from experts becomes complex and impractical in such a case [30, 54, 80]. One way to ensure that CPTs do not get too large is to invent hidden nodes [30]. Hidden nodes can make a network more compact as well as reduce the complexity by keeping the CPTs fairly small. An example of a hidden node in the rhino poaching network is *Accessibility*. The hidden nodes in the rhino poaching network were included during the expert workshop.

4.3.2 Routine Activity Theory

Another important concept is one that concerns the criminal aspect of the rhino poaching problem. Routine Activity Theory (RAT) is one of the main theories of criminology. It was developed by Cohen and Felson [81] who chose to focus on the conditions under which a crime will occur rather than focusing on the profile of the offender. RAT states that three elements need to be in place for a crime to occur: (1) there has to be a criminal or offender present, (2) there has to be a target or victim present, and (3) there has to be the absence of a capable guardian. Figure 4.1 illustrates the concept.

Figure 4.1. Routine activity theory

In the rhino poaching problem the offender or criminal is the poacher, the target or victim is the rhino, and the guardian is the ranger. The poacher will not be present (to poach a rhino) if the rhino is not present. The poacher will also not be present if the ranger is present, unless there is corruption involved, in which case it will not matter if the ranger is present or not. The ranger might be on leave or ill-disciplined in the sense that he leaks information, whether intentionally or by accident, thereby giving the poacher an advantage. The influence of this theory on the problem became evident after the expert workshop when subgroups in the network were identified. The subgroups that were classified based on RAT simplified future expert elicitation greatly, as will be discussed in the following chapters.

4.3.3 Transdisciplinary research

A variety of stakeholders across a wide range of fields need to be consulted and their responses integrated to create an understanding of the complex human-environment system [37]. This in itself is not sufficient to solve the problem. According to Düspohl *et al.* [37], there has to be "buy in" from the stakeholders; they have to support the study and be willing to implement the necessary strategies.

One way to consolidate the efforts of various stakeholders is to use the transdisciplinary (TD) research approach. TD research, also known as "participatory integrated assessment" or "post-normal research", is collaboration between experts by the generation of three types of knowledge [37]. The first type of knowledge ("system knowledge") is the knowledge that experts have of the system. The second type of knowledge ("target knowledge") is the knowledge that needs to be generated and integrated concerning the different perspectives, opinions, and objectives of the stakeholders. The third type of knowledge ("transformation knowledge") is the knowledge of how to achieve the common objectives.

According to Düspohl *et al.* [37] a fairly simple modelling approach is needed to support TD research which consists of three elements. The modelling approach should be able to "(1) represent and integrate knowledge from diverse disciplines and spheres, (2) explicitly support the inclusion of stakeholder knowledge and perspectives, and (3) take into account the uncertainty of knowledge" [37].

According to Kumar *et al.* [30] there are five main reasons why integrative models are developed. These reasons are (1) prediction, (2) forecasting, (3) decision making and management, (4) social learning, and (5) to develop an understanding of the system. The aim of this research study lies with prediction, decision making and management, and developing an understanding of the system. Locations for a possible next poaching incident is predicted (prediction), decisions are made concerning the allocation and management of resources such as rangers (decision making and management), and an understanding is developed of the complex rhino poaching problem (developing an understanding of the system).

BNs are well suited to address TD research problems due to their flexible nature, as well as their strong reliance on integration of stakeholder input. Problems that are assessed with a TD approach usually do not have enough data to automatically generate the network structure from the data [37]. In such cases, the network structure and the definition of the nodes and states are created in conjunction with the experts, or from literature, as is the case with this study.

4.4 DETAILED METHODOLOGY

This section elaborates on the research designs that were identified in Section 4.2. That, together with the concepts that were explained in Section 4.3 will aid the reader in understanding the methodology of the study. The study explores whether it is possible to predict the probability of a rhino poaching event in a specified geographical area and time epoch. Figure 4.2 represents a flowchart of the methodology that is followed in addressing the rhino poaching problem in this study.

The process illustrated in Figure 4.2 is akin to the "spiral approach" for constructing and evaluating the model. Although not appearing to be a spiral in the figure, it resembles the approach originally created by Boehm [82] and altered by Korb [75]. The blue blocks in Figure 4.2 correspond to parts of the study where expert knowledge is used, and the pink ovals are instances in the study where significant changes were made to a model. The order of the ovals represents the evolution of the models from a first-order (prior perspective) model to one that was evaluated and validated by experts. The green blocks denote instances where other "background" processes were executed. Green blocks always point to blue blocks or pink ovals, denoting extra knowledge that was needed or extra work that had to be performed in order to proceed in the flowchart.

Starting at the block named "Bayesian network work group" and following the arrows all the way to "Evaluate, validate" yields the process flow of the study and the thesis. Each of the blocks and ovals

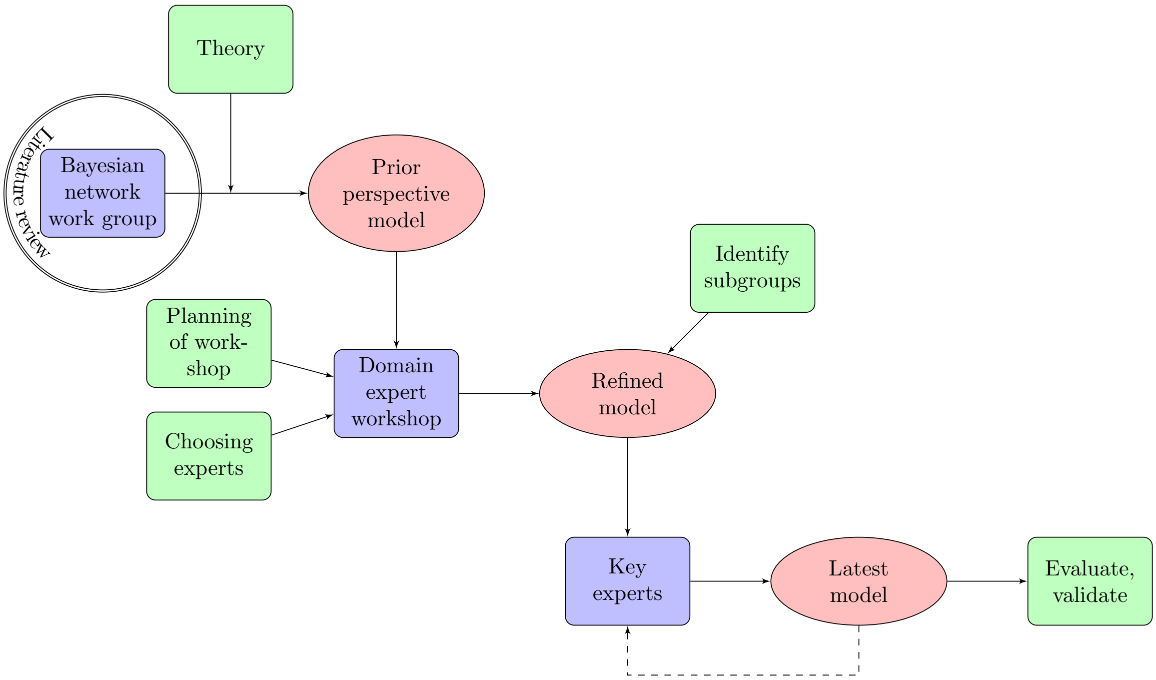

Figure 4.2. Flowchart of the study's methodology

will now be described and linked to the chapters and sections where they occur in the thesis. Figures 4.2 - 4.6 accompany the various sections.

Starting from Scratch

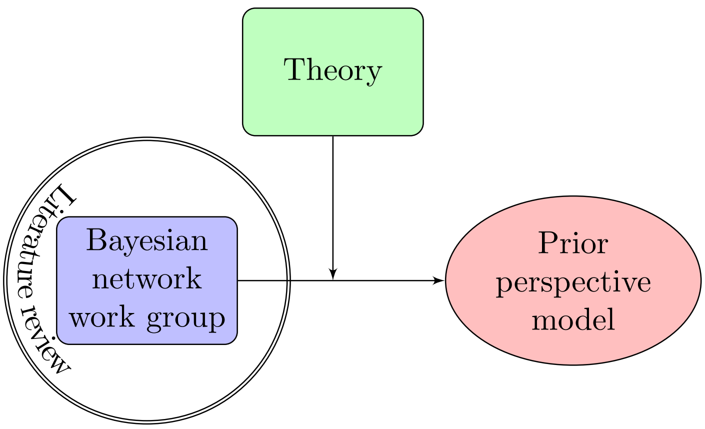

Figure 4.3. Starting from scratch

Referring to Figure 4.3, a prior perspective BN was developed by a small group of BN experts using literature (theory) and information obtained from informal discussions. In order to develop such a model, knowledge and theory was needed concerning BNs, as well as knowledge of the rhino poaching problem. This first-order BN was referred to as the “prior perspective model”.

The group of BN experts developing the prior perspective BN is an example of the survey-based research design that was described in Section 4.2.1. The development of the BN in terms of structure and contents is an example of statistical modelling which was described in Section 4.2.3. An in-depth discussion of the prior perspective model is described in Chapter 5.

The Expert Workshop

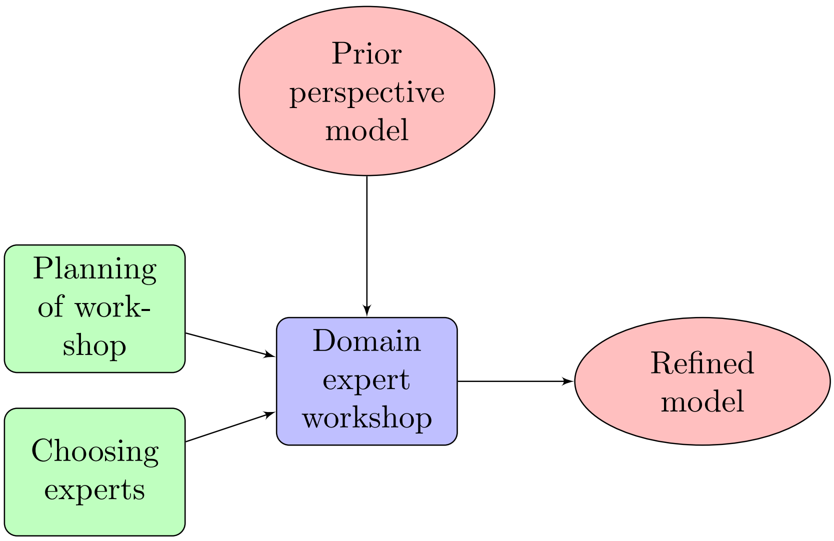

Figure 4.4. The expert workshop

Owing to the unavailability of sufficient rhino poaching data at that stage of model development, it was decided to use an expert knowledge approach to populate the CPTs of the prior perspective BN instead of taking a data-driven approach. The expert workshop was planned and the experts were chosen, as can be seen in Figure 4.4. The prior perspective BN was presented at the expert workshop for input and evaluation. After the expert workshop was concluded, a refined model was obtained that is a truer reflection of reality than the prior perspective BN.

The expert workshop was held for the evaluation of the prior perspective BN and is an example of the survey-based research design described in Section 4.2.1. Planning the workshop and choosing the experts, as well as arriving at a refined model, can be seen as examples of interdisciplinary research (as described in Section 4.2.4) as it required a study of different disciplines, by different disciplines. Arriving at the refined model is an example of statistical modelling as was discussed in Section 4.2.3. The workshop (planning and execution) and the refined model is explained in detail in Chapter 6.

The Latest Model

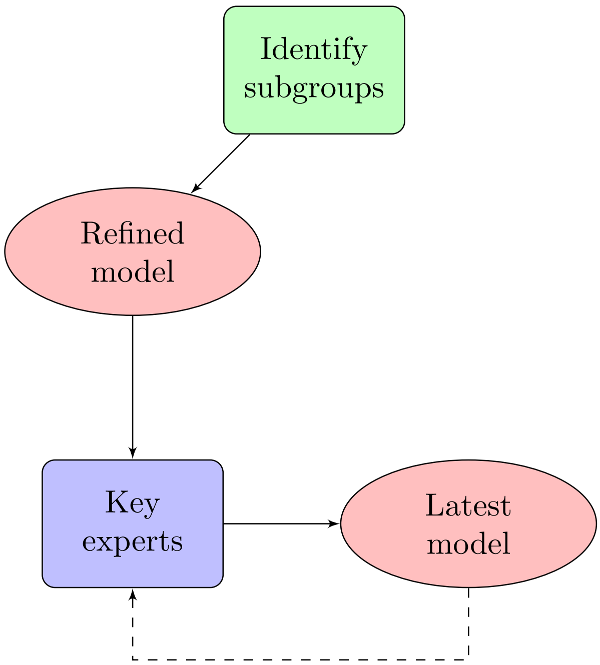

Figure 4.5. The New Model

After the workshop yielded the refined model, various subgroups were identified using the RAT principle. These subgroups made it easy to select key experts who could not attend the workshop. Meetings were set up with these key experts and each time the model was altered slightly, hence the feedback loop in Figure 4.5.

Communicating with the key experts is another example of the survey-based research design discussed in Section 4.2.1. Obtaining the latest model is examples of both survey-based research and statistical modelling (Section 4.2.3). Identifying the subgroups and eliciting knowledge from key experts is described in Chapter 5, and the latest model is explained in Chapter 6.

Evaluating the model

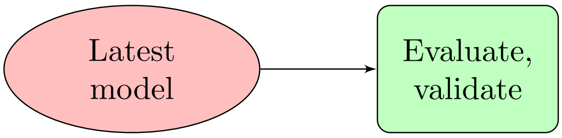

Figure 4.6. Evaluating the model

The last link in the process flowchart (Figure 4.6) is to evaluate and validate the latest model. The latest model was not trained on poaching data, but was rather parameterised using expert knowledge. Fellow researchers working on the broader rhino poaching project developed a BN from data (*data model*) based on this latest model. The latest model discussed in this thesis will henceforth be called the *expert model*. Evaluating the model requires analysing the secondary data as described in Section 4.2.2 and the evaluation and validation of the latest model is described in Chapter 8.

4.4.1 Research instruments

The BN was developed from narratives and a causal diagram was used in the workshop to obtain data in the form of knowledge from the experts. The output of the BN is a heat map that highlights possible poaching locations. This heat map is recursive and iterative and is also a research instrument. The heat map is recursive in the sense that the model is repeatedly applied to the problem of rhino poaching until a stopping condition is met, and it is iterative in the sense that the model will be executed and updated continuously as new data inputs are added.

Communication was established with various experts in the Kruger National Park (KNP) and data were obtained from them. In certain cases the author had to travel from Pretoria to the KNP (roughly 450 kilometres) to obtain the data on a storage device as it was too sensitive to share over public channels.

4.4.2 Expert knowledge and GIS data

The poaching record data are received in the form of a Microsoft Excel® (from here on referred to as "Excel") workbook. Various experts in the KNP captured the data. Rangers usually radio in new information and it is then recorded on paper and later typed into Excel by the GIS expert. Various hardware and software applications are currently being used to make the recording of the data easier and more accurate.

The first step is to prepare the Excel workbook for input to Matlab® (from here on referred to as "Matlab"). The original workbook is copied and then stored for safekeeping while the copy is used for processing. The Excel workbooks contain sheets with different variables, but the main variables are the poaching date (which is usually calculated between the date the carcass was found and the date of the necropsy) and the coordinates of where the carcass was found. All the entries are scanned for missing coordinates, as coordinates are among the most important elements of the data. The entries are also checked for inconsistency concerning format. For instance, is the column containing the poaching dates formatted in the same way? If the dates are not formatted the same it could cause significant problems when importing it into Matlab.

Some columns contain text, such as the region or property where the poaching attack occurred. Spelling or typing errors can also cause problems when importing, such as when the operator types "Stolsnek" instead of "Stoltznek", or "Nwanetsi" instead of "N'wanetsi", *etcetera*. Another error is when the operator enters the coordinates incorrectly. The coordinates are plotted in Matlab, and if there is a coordinate that is visibly in the "wrong" place, it is looked up in the Excel sheet, and more often than not it is due to a typing error. For instance, all the coordinates will start with -24 and 31, but the visibly wrong coordinate pair will start with 24 and 31. The error here is clearly that a "-" has been omitted.

Entries that are ignored are natural deaths of rhinos (*i.e.* not killed by poachers), as well as black rhinos. As mentioned, the scope of this thesis only includes white rhinos, as white and black rhinos have different habitats and behaviour. Also, white rhinos vastly outnumber black rhinos in the KNP. Lastly, any columns are removed containing variables that do not form part of the study. These variables include, but are not limited to, age, sex, or how the rhino was killed. Variables such as estimated date of death, weather, coordinates, and property are most important. Most of the variables in the model can be calculated from the estimated date of death or the coordinates of where the carcass was found.

All the workbooks are then merged into one and a check is performed to find duplicates. All duplicates are removed from the workbook. After the data has been cleaned and prepared, it is imported to Matlab. Importing an Excel file to Matlab is as easy as double clicking on the filename in the folder pane. Matlab then creates cell structures, which contain the data. These cell structures can be saved as variables, and loaded every time the code is run. When working in Matlab, the Mapping Toolbox® is invaluable. It lets the researcher import and work with shape files (digital maps). The data points of the poaching attacks can directly be overlaid on a map of the area.

The rhino poaching data are used for training the alternative data-driven model as well as possibly testing both models in order to compare them. The other source of "data" that was used in this study is expert knowledge. Expert knowledge differs from captured data in the sense that expert knowledge is tacit knowledge, versus the explicit knowledge of data. Data capture details of events that have already occurred, while expert knowledge is more subjective and flexible.

4.5 LIMITATIONS

All the data and perceptions of rhino poaching are driven by specific research questions that the original researcher(s) had. This biases the data collection in a way, as different perspectives are not considered while various stakeholders have their own opinions of the problem. During the study we came to the conclusion that rhinos are not the victims or targets, but the commodity in the problem. This changes the entire problem in terms of how data are captured and how the problem, in general, is viewed. The fact that rhinos are viewed as the victims or targets of the problem necessitates agencies to gather data whereby rhino poaching events are the main focus of the data. The date, location, and other pieces of information regarding the carcass are gathered, but none regarding, for example, the socio-economic information.

During the past few years many policy changes have been made in the KNP. Each time a major policy change is made it affects the way in which the poachers react. For instance, rangers always had somewhat fixed patrol routes, and when the persons in control realised that poachers prefer to poach during the full moon, they sent the rangers to those areas every full moon. The poachers then had to adapt their behaviour to no longer be predictable. Every time such a change occurs it changes the pattern of the data. It is unknown when which policy changes occurred, hence it is not known on which intervals of data to train which variables. Some variables might have been important two years ago (say, moon phase), but it is no longer such an important variable. This greatly complicates the training and testing of the models. Given all of these limitations, the next few chapters will be used to shift the rhino poaching problem to a space where it can be understood and where researchers can reason about it.

4.6 ETHICS

Poachers are killing rhinos, and the armed forces in turn are killing or arresting the poachers. This work, however, only serves to comment on the impact of various trends on the reduction of rhino numbers. There was no contact with poachers nor rhinos from the author, and all the data used in this study are secondary and was collected and captured by rangers or other individuals in the KNP. The information that was verbally obtained was done so with the consent of all the experts involved. Ethical clearance was obtained from the University of Pretoria as well as the Council for Scientific and Industrial Research (CSIR) to conduct an expert workshop whereby the knowledge of the experts was elicited. The results from the workshop were recorded and dictated while keeping the experts' identities anonymous.

4.7 CONCLUSION

This chapter presents the methodology used in the study to shift the rhino poaching problem to a space where researchers can begin to address it in order to arrive at a conclusion on how to mitigate the problem. It became clear that the rhino poaching problem has multiple dimensions and that a single approach would not suffice in addressing it. Different methodologies had to be investigated and used to address the different dimensions of the problem. Transdisciplinary research allows for collaboration with a wide variety of experts as well as to facilitate the co-creation of the BN, thereby shifting the problem closer to a reasonable resolution.

The study was started from scratch - a small group of BN experts developed a prior perspective BN. This prior perspective BN was then evaluated at an expert workshop where it was altered until the experts all agreed with the structure and semantics of the network. After the workshop several subgroups were identified in the structure of the model, which made it easier to further elicit knowledge from experts. Key experts who could not attend the workshop were identified and meetings were arranged. After each meeting with a key expert the model changed slightly until the model structure started to converge. No model will ever be precise or “finished”, thus it was decided to use the structure convergence as the stopping criterion.

CHAPTER 5 DEVELOPING THE BAYESIAN NETWORK

5.1 INTRODUCTION

Rhino poaching is a highly complex problem and at the start of this project very little was known about the problem. Most of the information came from field rangers and criminal investigators working on rhino poaching cases. There were no studies highlighting the possible causes and patterns leading to rhino poaching. There was also no scientific literature to consult pertaining to rhino poaching. The author wanted to contribute to the anti-poaching efforts in a way that had not been addressed previously, namely that of gaining an improved understanding of the poaching problem, predicting poaching hotspots, and alerting the necessary parties to probable times at which poaching events could occur. A model was developed based on domain expert knowledge, without using prior work or knowledge. As mentioned earlier in this document, an initial causal Bayesian Network (BN) model was developed to model the rhino poaching problem.

This chapter (as well as the next) is devoted to the development of the BN model from conception to its current working form. It is worth noting that the research project progressed from a state where the problem was unknown and highly complex, to a state where the problem can be addressed in a systematic manner by different groups of researchers and stakeholders. This was achieved by developing a simplified structured view of the problem, while capturing its most influential factors or components.

5.1.1 The process

A first-order or prior perspective BN model was developed by a small group of BN experts and was presented at a workshop consisting of domain experts. The workshop was planned, organised and facilitated by the author. The author also chose which participants to invite. This prior perspective model was used instead of the more conventional "clean slate" approach where the experts have to come up with their own model during the workshop. The prior perspective model was developed before the workshop in order to save time and to reduce the risk of expert unavailability. This will be explained further in Section 5.3.

Chapter 4 mentioned and discussed the "spiral approach" by using a flowchart of the study. The prior 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 almost uniform agreement.

5.2 DEVELOPING THE FIRST ITERATION MODEL

5.2.1 The original model

A small group of BN experts developed the original model shown in Figure 5.1. These experts included individuals from the Council for Scientific and Industrial Research (CSIR) Pretoria (including the author) as well as colleagues from Thales Research and Technology in Delft, The Netherlands. Initially there was limited information available and relationships between the Kruger National Park (KNP) and researchers had to be cultivated. This was a challenge for the model. The BN experts had to use information obtained from informal personal communications, their knowledge of BNs, literature, as well as insight to develop a logical model.

The model contains variables pertaining to the rhino poaching problem all leading to *Poaching_event* which is the target variable of the model. It was decided to divide the landscape into a grid of cells, and the variable *Cell_location* then points to the specific cell that is under analysis. The variable *Cell_size* then specifies the size of the grid (5×5 kilometres, 10×10 kilometres, *etcetera*). A prediction

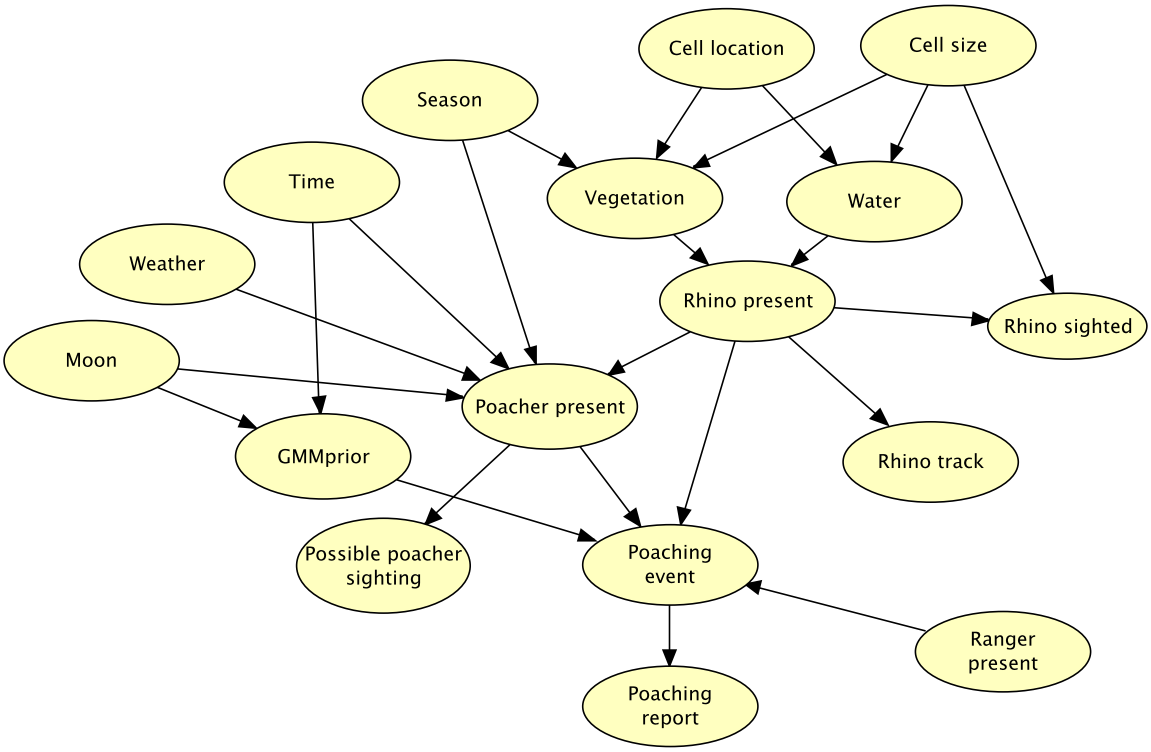

Figure 5.1. Original BN developed by a small group of BN experts

is made for a specific cell, thus it is not necessary to specify it in the network itself. The variable *Cell_location* would also have had hundreds or thousands of states, depending on the grid size. There would have been the same problem with *Cell_size*, as each possible cell size combination would have to be accounted for in the list of states. At first *Cell_location* and *Cell_size* were treated as random variables, but they are always observed, hence they were excluded from future models.

The *GMMprior* variable encapsulates the effect of previous poaching attacks. It acts as a timeline of poaching records and relates back to the model as a prior probability. The hypothesis behind this is that poaching attacks tend to occur at the same locations. Thus, if a cluster of historic poaching attacks occurred at some location X, then the cell containing location X will have a higher prior probability of being home to a future poaching location, without the system knowing anything else about that cell. The variable is called *GMMprior* because the first approach to calculating these probabilities was to use Gaussian Mixture Models (GMMs) to attempt to cluster the different variables in space.

5.2.2 The prior perspective approach

The model was studied amongst the experts in the group and a few changes were made. The model presented in Figure 5.2 is the prior perspective model and is the model that was presented to domain experts at the workshop.

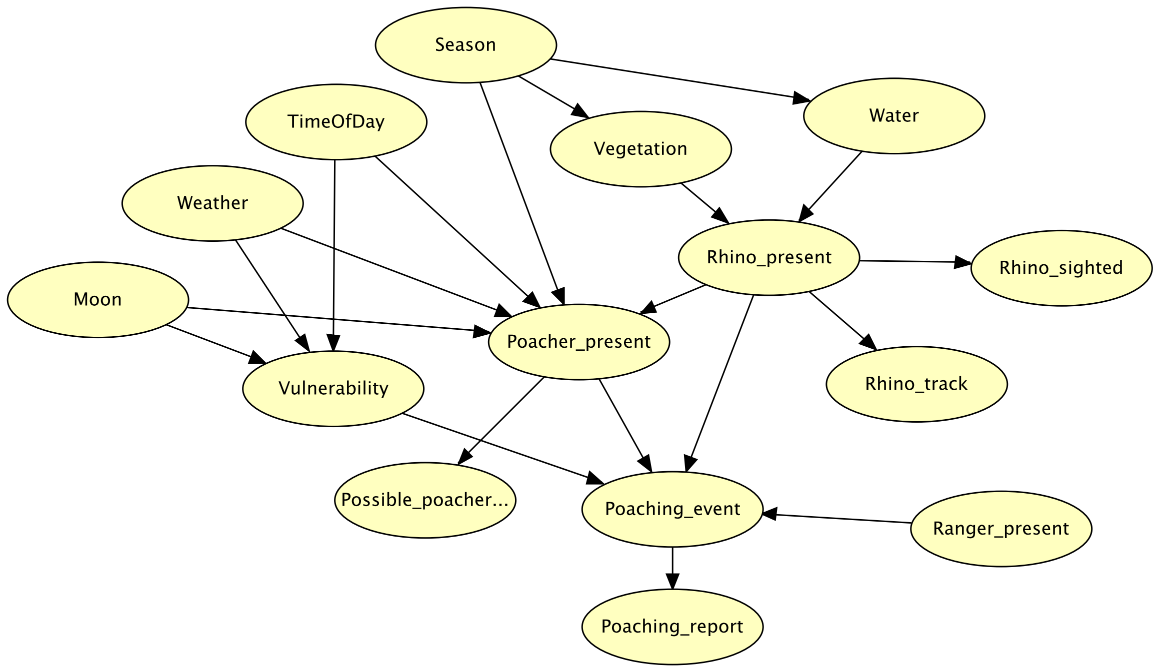

Figure 5.2. Prior perspective model

The main difference between the original model and the prior perspective model is the removal of the *Cell_location* and *Cell_size* variables. As mentioned above, the cell location and cell size is implied in the model, and are always fully observed. Hence, they are not represented by random variables. The *GMMprior*'s name was also changed to *Vulnerability*, as it was felt that this more accurately reflected the meaning of the variable. The approach to calculating the probabilities is also not limited to only using GMMs.

5.3 BEFORE THE EXPERT WORKSHOP

Traditionally, workshops are held according to the "clean slate" approach. This entails that all the participants (including the organiser and facilitator) come to the workshop with a blank slate, and the participants are tasked to develop a model from scratch. The workshop participants then design and develop their own model. For this workshop, however, a first-iteration model was developed beforehand and given to the experts for evaluation and validation. The reason for the workshop was to evaluate the prior perspective model and to elicit expert knowledge from individuals working with rhinos or working in game reserves. The idea was that the experts evaluate the structure of the model and provide guidance in populating the Conditional Probability Tables (CPTs) with probabilities. Real-life data for the rhino poaching problem are scarce, incomplete, and very sensitive, as is the case with most environmental problems [45, 47, 83, 84].

It is relatively easy to get a group of people together for a meeting if the duration of the meeting is an hour or two. Organising a workshop that stretches over (at least) two days with experts from across the province poses problems, such as time and availability. The rhino poaching problem is especially prevalent in the KNP which is roughly 450 kilometre from South Africa's administrative capital, Pretoria (where the researchers are based and the workshop was held). Many experts who have the necessary knowledge have to travel from the KNP to Pretoria for the workshop and it is expensive to travel that distance both by car and by air.

The rhino poaching problem requires constant attention and reaction, thus some experts have to return to the KNP at a moment's notice. Members of the criminal investigation unit are on call at all times and might not even make it to workshops or meetings due to urgent work matters. Personal communications with a member of the Green Scorpions (Environmental Management Inspectors in South Africa) revealed that some of their members do not know their own whereabouts more than two hours in advance. When taking all these factors into account it makes sense to develop a first iteration BN. It is quicker and easier to make changes to an existing model than to build a new model from scratch.

The prior perspective model was developed to allow for the time and availability constraints of the experts. A possible shortcoming of having a workshop with a pre-defined model, is the possibility of biasing domain experts towards a particular view of the problem, but Finlay [85] points out that the key to the introduced reflexivity is to be sensitive to how this is done. This has to be traded off against the savings in man-hours and other considerations when assembling such a group of domain experts.

5.3.1 Planning the workshop

A workshop can be thought of as a scientific experiment: the organiser sets up the experiment and executes it. The organiser decides who will be the participants in the experiment, and also decides on focus questions to direct the flow of communication. It is important to decide which areas need to be covered, and by whom. Different stakeholders or experts will have different opinions and areas of expertise. It is important to have a variety of stakeholders present to ensure an objective workshop with objective outcomes. There are often experts who can supply researchers with knowledge where data are either unreliable or unobtainable [44]. According to Fisher *et al.* [44], a notable increase in the use of expert elicitation in ecological applications is observed due to the fact that data are not easily obtainable.

Another important factor to consider is the type of elicitation style that will be used. There are many different elicitation styles, which include face-to-face type interviews, surveys, and teleconferences. According to O'Leary *et al.* [47], the elicitation technique used is contingent on the type of expert as each expert has different levels of statistical knowledge and mapping skills. The style of this workshop was face-to-face as it is much easier to communicate exact thoughts and ideas when speaking to someone directly [86]. If an expert is asked a question face-to-face, it is easy to establish whether or not the expert has understood the question. The expert can also ask specific questions to make sure he understands, or ask additional questions. It is also easier to draw pictures and ideas on a piece of paper or on a whiteboard if something is unclear, instead of trying to explain it via Skype or email. The elicitor can also ask questions specifically formulated to suit the expert's skill level pertaining to probabilities [44, 47].

5.3.2 Choosing the participants

A very important aspect of expert elicitation is choosing the experts [47]. An individual can be an expert due to their education, training, experience, or a combination of these factors. There are no exact criteria by which to establish when an individual is an expert.

In the book by Esbjörn-Hargens and Zimmerman [87], the case is made for integral ecology. Integral ecology draws together different disciplines and answers the call for greater multidisciplinary as well as transdisciplinary collaboration. This is the approach that is followed in this study as the rhino poaching problem stretches across many more areas than just rhino ecology. No single team or department can solve the rhino poaching problem, collaboration is key. The experts for the workshop were chosen to represent the quadrants of the mental model explained by the authors [87]. There are four areas of expertise that can be used for enquiry, analysis, investigation, and assessment. The context for the mental model in this case is rhino poaching. The current reality is that more than a thousand rhinos are poached per year in South Africa.

The four areas of the mental model that can be covered are the psychological and phenomenological inquiry of the problem, the behavioural and physiological analysis of the problem, the ecological and social assessments, as well as the cultural and worldview investigations of the problem. Not all of these areas are represented fully in the workshop (psychological and phenomenological inquiry and cultural and worldview investigations) as they fall outside of the scope of this study. The behavioural and physiological analysis of the problem and the ecological and social assessments are represented by experts in the workshop, but all four areas work together to place the work in context within the literature.

The type of expert needed for psychological and phenomenological inquiry is an individual who can give a first-person perspective on the psychology behind rhino poaching. The type of expert needed for the behavioural and physiological analysis is an individual who collects and works with the rhino data such as an ecologist, game ranger, analyst, or someone in the operations room of a game reserve.

The type of expert needed for the ecological and social assessments is the same as for the behavioural and physiological analysis, but can also include government officials concerned with environmental problems, as well as safety and security (political and legal). Members of the army deployed in the KNP would cover the political, environmental, as well as legal aspects. Academics at universities who study the political, educational, legal, and environmental factors, as well as economists and members of the World Wildlife Fund (WWF) could also make valuable contributions.

The type of expert needed for the cultural and worldview investigations is an individual who has a background in philosophical, ethical or religious factors pertaining to the problem. Possible experts could be academics who focus on these factors, as well as sangomas (traditional African healers) or poachers themselves. Getting hold of poachers or sangomas will prove to be a fairly difficult task, and an alternative to this is to do a thorough study of the literature pertaining to the cultural perspectives surrounding the problem. The core of this quadrant (in this specific case) is the ethics surrounding the problem – the chain of consumers and distributors who use local and other organisations for the work by hiding behind needs such as clothes, food, and political objectives.

After choosing the experts according to these criteria they were contacted and the problem was described to them, as well as the reasons for the workshop. The roles of the experts were made clear to them, although specific roles were not assigned to each expert beforehand as it was unclear as to what each expert's knowledge level was.

5.4 DURING THE EXPERT WORKSHOP

During the introduction of the workshop, the participants were asked to make a list of their expectations, not only of the workshop, but also of how the workshop could add value to their own work. The experts wanted to reach as much consensus as possible, bring a different perspective to the rhino poaching problem, and obtain a workable solution that is manageable on all levels. They also wanted to have a common understanding of a BN, to know what it can be used for, and how to apply it in the real world. The experts also wanted a realistic model which captures their inputs correctly in order for them to understand the patterns of poaching. It was thought that networking could assist in counter-poaching and that the model could even be expanded to other environmental applications of crime prevention.

During the workshop, the experts were briefly trained in the concepts and use of BNs so that they had a basic understanding of causal models and would be able to understand a BN model. Figure 5.2 illustrates the original prior perspective model that was showcased at the workshop. This model was presented for the first time in the paper by Koen *et al.* [25].

The model was given to the experts as a handout to take home and study. The handout also contained a breakdown of the different nodes, states, and why they were included. Key points the experts had to consider were (1) which variables and links they would remove or change, (2) what they would add to the model to make it more realistic, and (3) if the semantics of the model are correct.

On the second day of the workshop there were a lot of questions and comments about the network. The experts understood the BN and were excited to contribute to the discussion. A few of the suggestions for new variables (nodes) and links were implemented in the network, but had to be explained to the experts as it was unclear from the handout and their limited understanding of BNs.

There is no effective measure to quantify whether or not the prior perspective model is better than the "clean slate" approach. However, after presenting the model to experts and letting them think it over, the model they came up with was remarkably similar to the prior perspective model, in both structure and meaning. This could be attributed to "group think" (where one person in a group offers an opinion and the rest of the group agrees with him without suggesting alternate options), but the experts each had time to think about the model on their own.

Developing an independent BN before the start of the workshop offers researchers the opportunity to let experts evaluate and validate the model during the workshop. Having a pre-defined model ready can somewhat accelerate the process of eliciting knowledge. After the experts worked on evaluating the structure of the model, the next step was for them to populate the model with probabilities. This is discussed in the next section.

5.4.1 Eliciting probabilities from participants

Expert knowledge elicited from workshops can be used in two ways, namely to act as a prior distribution for the parameters, or to serve as a comparison against empirical data or empirical estimates. These starting estimates describe the present situation of the knowledge when acting as a prior distribution, and are updated as new information becomes available [44].

Breaking the problem down into a series of small questions that are well defined are beneficial for the researcher as well as the expert. "Model decomposition in this way helps identify components that the expert may find easier to relate to, but which may also be independently elicited" [44]. Even if the expert knowledge elicited is not truly accurate, it still presents the researcher with an interesting template to compare to the empirical data, or subsequent expert elicitation sessions [47].

There exist several biases that the expert and the facilitator are faced with, such as displacement (the expert over/underestimates the expected value), variability (the expert underestimates the certainty of this response), motivational bias (the expert's opinions are influenced for some reason), and cognitive bias [47, 78]. Cognitive biases include availability bias, anchoring, adjustment, representativeness, overconfidence, conjunction fallacy/coherence bias, and hindsight bias [44, 47]. *Availability bias* arises when the expert recalls only recent or important events. *Anchoring* is when the expert adjusts all of the probabilities and values to fit his best guess. There are methods that can be deployed to curb these biases, but some biases are, however, not as easy to control. The expert might know about previously calculated estimates (which were either incorrect or repeated numerous time in the literature) and that might influence his prior perspective or answer to the question. As Fisher *et al.* [44] states, it "creates a false sense of security" with regards to the expert's estimation of the parameter.

The easiest way of eliciting probabilities from experts is to ask them for a "baseline" and then to adjust this baseline up or down depending on the different states of the variables. Experts usually find it easier to adjust numbers up or down, rather than thinking of numbers from scratch [44].

Table 5.1 shows an example of what the Conditional Probability Table (CPT) could look like for *Rhino_present* if *Landscape_preference* and *Water* feeds into it as in Figure 5.3. This simple example illustrates the concept of eliciting expert knowledge which can be used in much more complex examples as will be discussed in the next chapter.

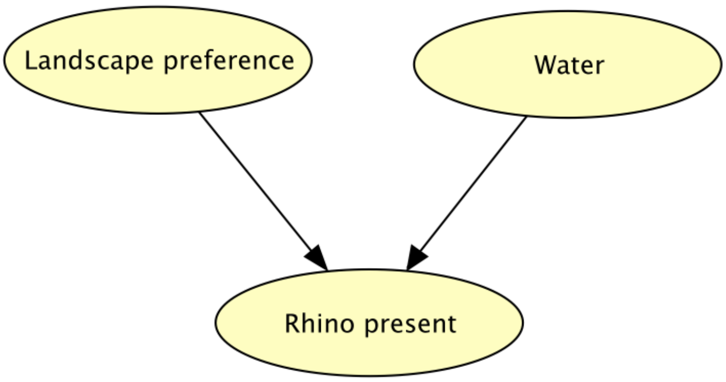

Figure 5.3. *Rhino_present* subnetwork

The expert is asked which of these variables feeding into *Rhino_present* is, according to him, the most important for establishing knowledge about the target variable (*Rhino_present*)? Assume the answer is *Landscape_preference*, then that node is moved to the top of the CPT as illustrated in Table 5.1. The rest of the nodes are then moved below this node, but in this case there is only one, *Water*. The

Table 5.1. *Rhino_present* example

Landscape_preference	Neutral		Avoid		Prefer	
Water	Not in cell	In cell	Not in cell	In cell	Not in cell	In cell
False	*	*	*	*	*	*
True	*	*	*	*	*	*

states of the target variable (*Rhino_present*) is then placed on the left hand side of the table, namely "False" and "True".

Next, the expert is asked which of the states of this important variable he feels most confident with to answer questions on. Assume the answer is "Prefer". The expert, with the help of the facilitator, then populates the "Prefer" section of the table with probabilities. The expert now only needs to fill in two probabilities: $P(Rhino\_present$ = "True"$|Landscape\_preference$ = "Prefer"$, Water$ = "Not in cell") and $P(Rhino\_present$ = "True"$|Landscape\_preference$ = "Prefer"$, Water$ = "In cell"). The corresponding values for $P(Rhino\_present$ = "False"$|Landscape\_preference$ = "Prefer"$, Water =$ {"Not in cell", "In cell"}) will just be $1 - P(Rhino\_present$ = "True"$|Landscape\_preference$ = "Prefer"$, Water =$ {"Not in cell", "In cell"}). Table 5.2 shows the table after the first probabilities pertaining to "Prefer" have been completed.

Table 5.2. *Rhino_present* example

Landscape_preference	Neutral		Avoid		Prefer	
Water	Not in cell	In cell	Not in cell	In cell	Not in cell	In cell
False	*	*	*	*	0.3000	0.1500
True	*	*	*	*	0.7000	0.8500

Next, the experts have to discuss what the weighting of the other states are in relation to "Prefer", which was deemed the most important. If "Prefer" has the most weight, what will the respective weights of "Avoid" and "Neutral" be? If the experts decided that "Avoid" would carry a 0.4000 weight, and that "Neutral" would carry a 0.6000 weight, the two states "Avoid" and "Neutral" can then be populated by using the probabilities of "Prefer" together with the weighting factor. The probability of a rhino being present given that *Landscape_preference* = "Avoid", and *Water* = "Not in cell", is simply $0.4000 \times 0.7000 = 0.2800$. The probability of a rhino not being present given these same state values will then be 1 - 0.2800 = 0.7200, and so forth. Table 5.3 illustrates the CPTs if the weighting scheme is used as above.

Table 5.3. *Rhino_present* example

Landscape_preference	Neutral		Avoid		Prefer	
Water	Not in cell	In cell	Not in cell	In cell	Not in cell	In cell
False	0.5800	0.4900	0.7200	0.6600	0.3000	0.1500
True	0.4200	0.5100	0.2800	0.3400	0.7000	0.8500

This method can be used when the experts are able to assign probabilities or uncertainties to scenarios, but also when the size of the CPT is too big to go through it state by state. A CPT of 2×6 was effectively reduced to a table of 1×2.

5.4.2 Evolution of the network

Figure 5.4 shows the model after it was first altered during the workshop. The model's structure evolved from the prior perspective model to a more detailed model. This section deals with the evolution of the network up until the end of the workshop.

The main differences between the prior perspective model and the refined model are the nodes that were added. *Festive_periods* were added as the experts stated that the number of poaching attacks seems to increase with a proximity to festivals and holidays. A variable in the same class is *Stages_of_the_month* as the proximity to payday is also said to be a good indicator of poaching attacks. The *Vulnerability* node was changed to *Historical_vulnerability* to more accurately reflect the fact that historical data is being used.

Other nodes added are *Crossing_points* (where do poachers enter and exit the park?), *Traffic* (does heavy traffic or no traffic make a difference to poachers?), *Distance_to_human_settlements* (the poachers can disappear in human settlements), *Vegetation_density* (if the poachers cannot enter densely populated areas they will not go there), *Soil_moisture* (an indicator of good grazing areas), *Fire* (rhinos love the soft tufts of grass that sprout two weeks after a fire), and *Grazing_availability* (measures whether or not a rhino will eat the grass at that location).

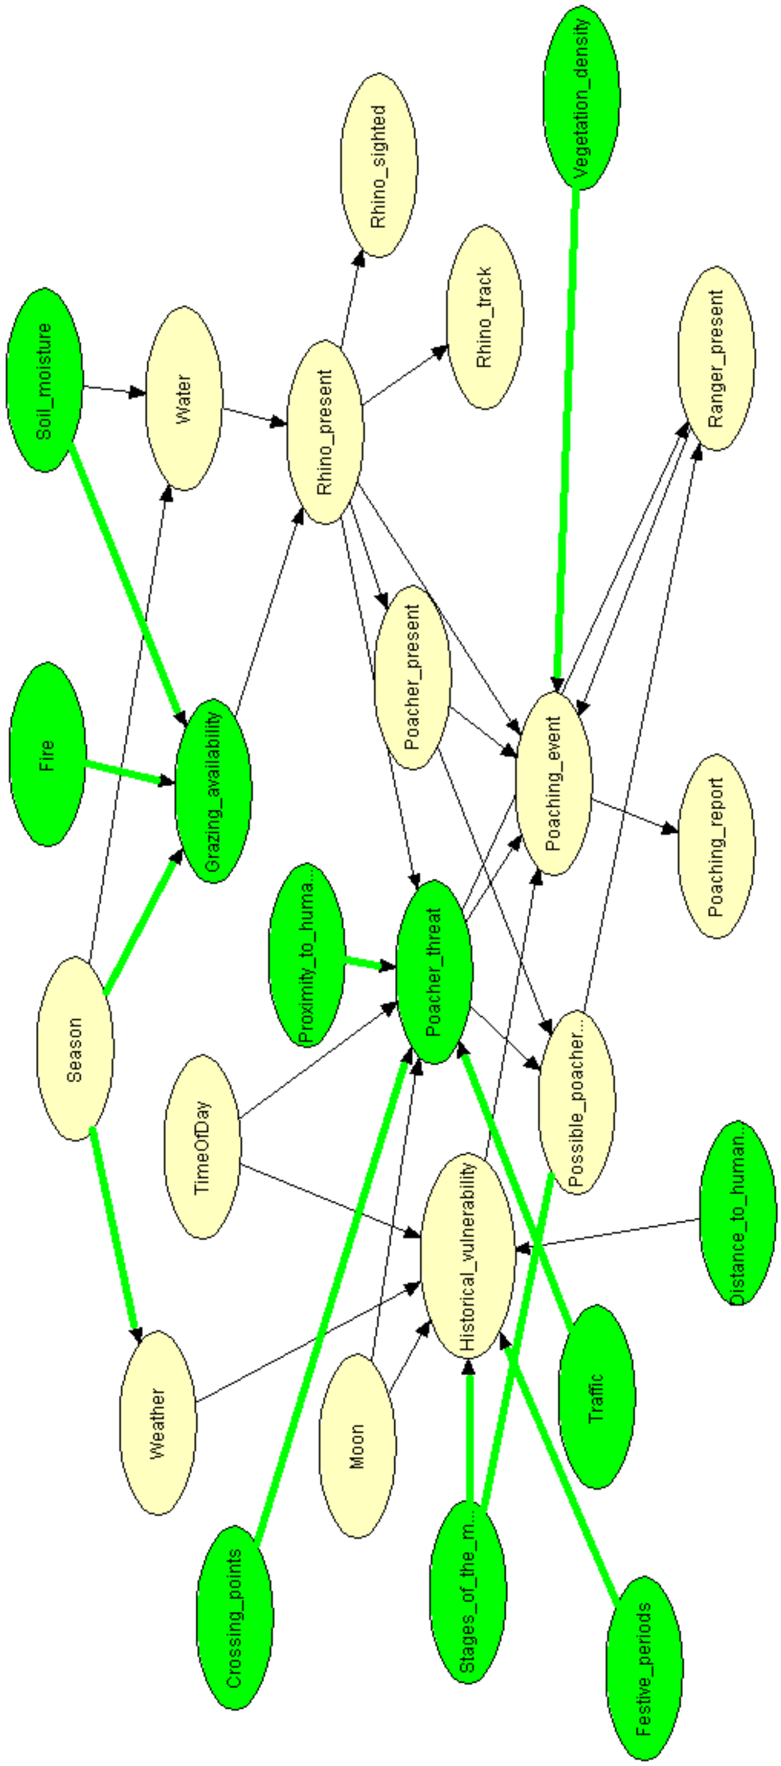

Figure 5.4. First model altered during the workshop

Figure 5.5 shows the final model at the end of the workshop. The number of changes that were made to Figure 5.4 to obtain the model at the end of the workshop were smaller than in the previous iteration. The *Traffic* and *Season* nodes were removed. *Traffic* was removed as it seemed too complicated to quantify effectively. *Season* was also removed due to the fact that some of its properties are already captured in other variables such as *Soil_moisture*. The *Vegetation_density* node's arc was moved from pointing to *Poaching_event*, to pointing to *Accessibility*. It was argued that the density of the vegetation would have a greater impact on a poacher's access into a certain area than on the probability of a poaching event directly.

Two nodes were created and others were renamed to clarify the meaning of the nodes. *Corruption_index* was added as it was argued that there is a measure of corruption present amongst the rangers and safeguarding personnel. A hidden node, *Accessibility*, was created to act as a proxy between the factors that influence a poacher's access into the park or into a certain area. Hidden nodes are used to keep the number of incoming edges into nodes in check. If *Accessibility* was not used, there would be five edges leading into *Poacher_present*, unlike the current three. The downside to using hidden nodes is that information could get lost.

The *Distance_to_human_settlements* node was changed to *Proximity_to_static_deterrents* as "static deterrents" refer to more than just settlements; it refers to such as camps, lodges, and ranger posts. *Crossing_points* was changed to *Distance_to_incursion_points* as this includes different types of entry and exit points. Lastly, *Ranger_present* was changed to *Active_deterrents* as it was argued that rangers are not the only entities able to deter poacher: other deterrents include alarms and sensors.

5.5 AFTER THE EXPERT WORKSHOP

After the workshop the information shared was processed along with the newly evaluated and altered model. While studying the model it became clear that the network could be divided into smaller subnetworks. The identified subnetworks were a rhino subnetwork, a poacher subnetwork, a ranger subnetwork, and a historical influence subnetwork.

Subnetworks make it easier to elicit expert knowledge in future. Referring back to Section 5.3.2, it is clear that not all experts are equipped to, or feel comfortable with, eliciting information pertaining to all quadrants of the mental model. By dividing the network into subnetworks, a specific subnetwork can be shown to the correct expert. The entire model could, for instance, be explained to a rhino ecologist, while only showing him the rhino subnetwork. He can then answer questions about the area of expertise that he is comfortable with. The various subnetworks of the current network are discussed in detail in the following chapter.

5.5.1 Further evolution of the network

As mentioned in Section 5.3, getting all the relevant experts together was a challenging task. The author arranged individual interviews with key experts who could not attend the aforementioned workshop. The subnetworks were shown to them and their inputs were garnered.

Figure 5.5 shows the latest model after the workshop, while Figure 5.6 illustrates the model after interviews with key experts. Most of the changes made between the two models were causal links between nodes that were moved. After deliberating with experts, certain aspects became clear. The edges linking the nodes *Festive_periods*, *Weather*, *Stages_of_the_month*, *Moon*, and *Time_of_day* to *Historical_vulnerability* were removed as it was argued that these variables were already implied in the historical data used for *Historical_vulnerability*. The edges from nodes *Festive_periods*, *Weather*, and *Stages_of_the_month* were now linked to *Active_deterrents*. *Moon* and *Time_of_day* were directly linked to *Poacher_present*. The *Season* node was also re-added (the *Rhino_present* subgroup was a continuous point of discussion during the development of this model), as was a *Vegetation_type* node to help identify the *Grazing_availability*.

The differences between the models in Figure 5.6 and Figure 5.7 are due to interviews held with GIS specialists in the KNP. The arc of *Weather* was changed to not flow into *Active_deterrents*, but rather into *Poacher_present*. The *Season*, *Fire*, and *Soil_moisture* nodes were removed. This was done at the recommendation of an expert who argued that the complexity of the *Rhino_present* subgroup will just continue to explode if it is not kept simple. It was decided that grazing and vegetation is beyond the scope of this study since it will require a separate investigation into those areas. The expert's advice was to keep the vegetation as simple as possible, thus only using a single node called *Vegetation_type*.

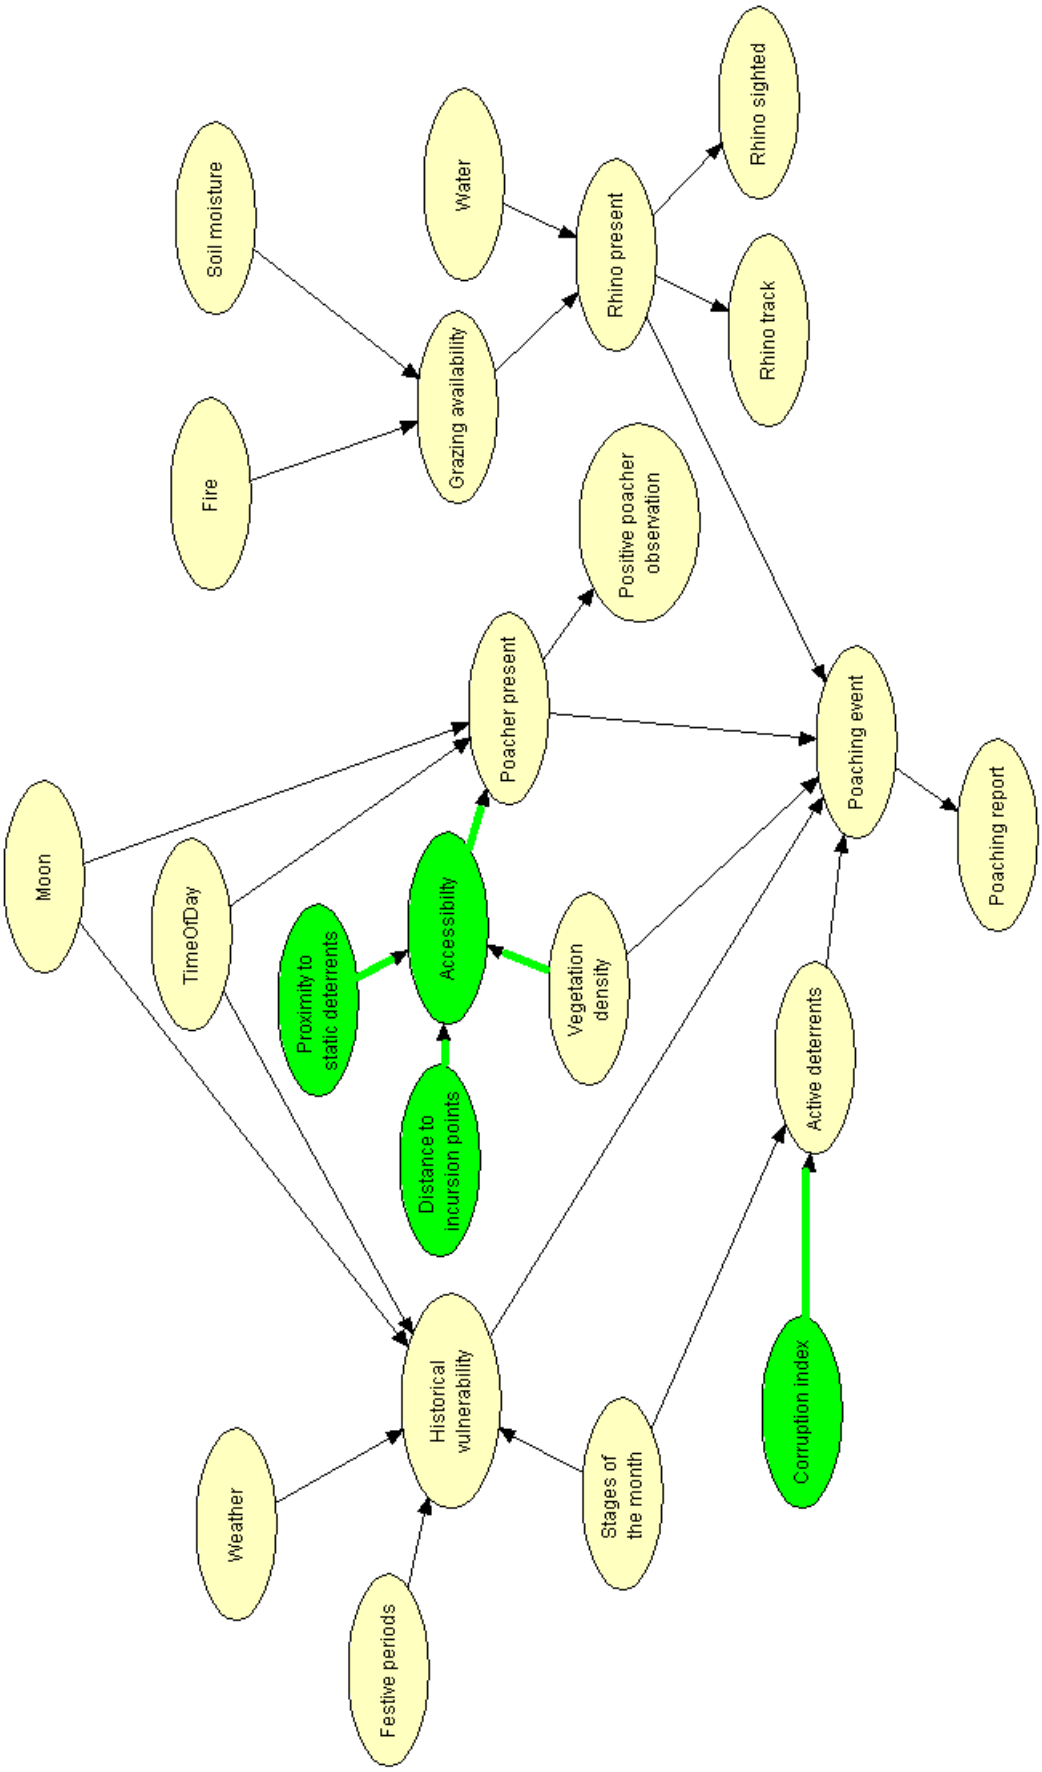

Figure 5.5. The refined model

New nodes that were added are *Slope*, *Zone*, and *Historical_rhino_presence*. *Slope* was added because white rhinos prefer undulating plains and avoid areas that are too rocky and steep. *Zone* was added to toggle a switch between the different protection zones. There are three protection zones in the KNP, each with its own distribution for rhinos. *Historical_rhino_presence* was added to incorporate historical rhino presence data such as sightings and census data. *Historical_rhino_presence* is linked to *Zone*, as that is what the switch is for.

The changes between Figure 5.7 and the latest model (Figure 5.8) were minor as the models now seemed to converge. The *Slope* variable was removed, and *Vegetation_type* was changed to include both the slope and the vegetation type into a new variable called *Landscape_preference*. This variable states whether or not a rhino will be at a certain location, whether it be muddy pools, good vegetation, or cool grass to lie on. The node *Moon* was also changed to become *Moon_illumination_percentage*, because how much can be seen by the light of the moon is more important than what phase the moon is in. The biggest change to the network, however, is the removal of the *Historical_vulnerability* node. The reason for this is that *Historical_vulnerability* contains data from historical poaching records, and it was decided to not use data in the model. We wanted to see what the effects of a purely expert-driven model is.

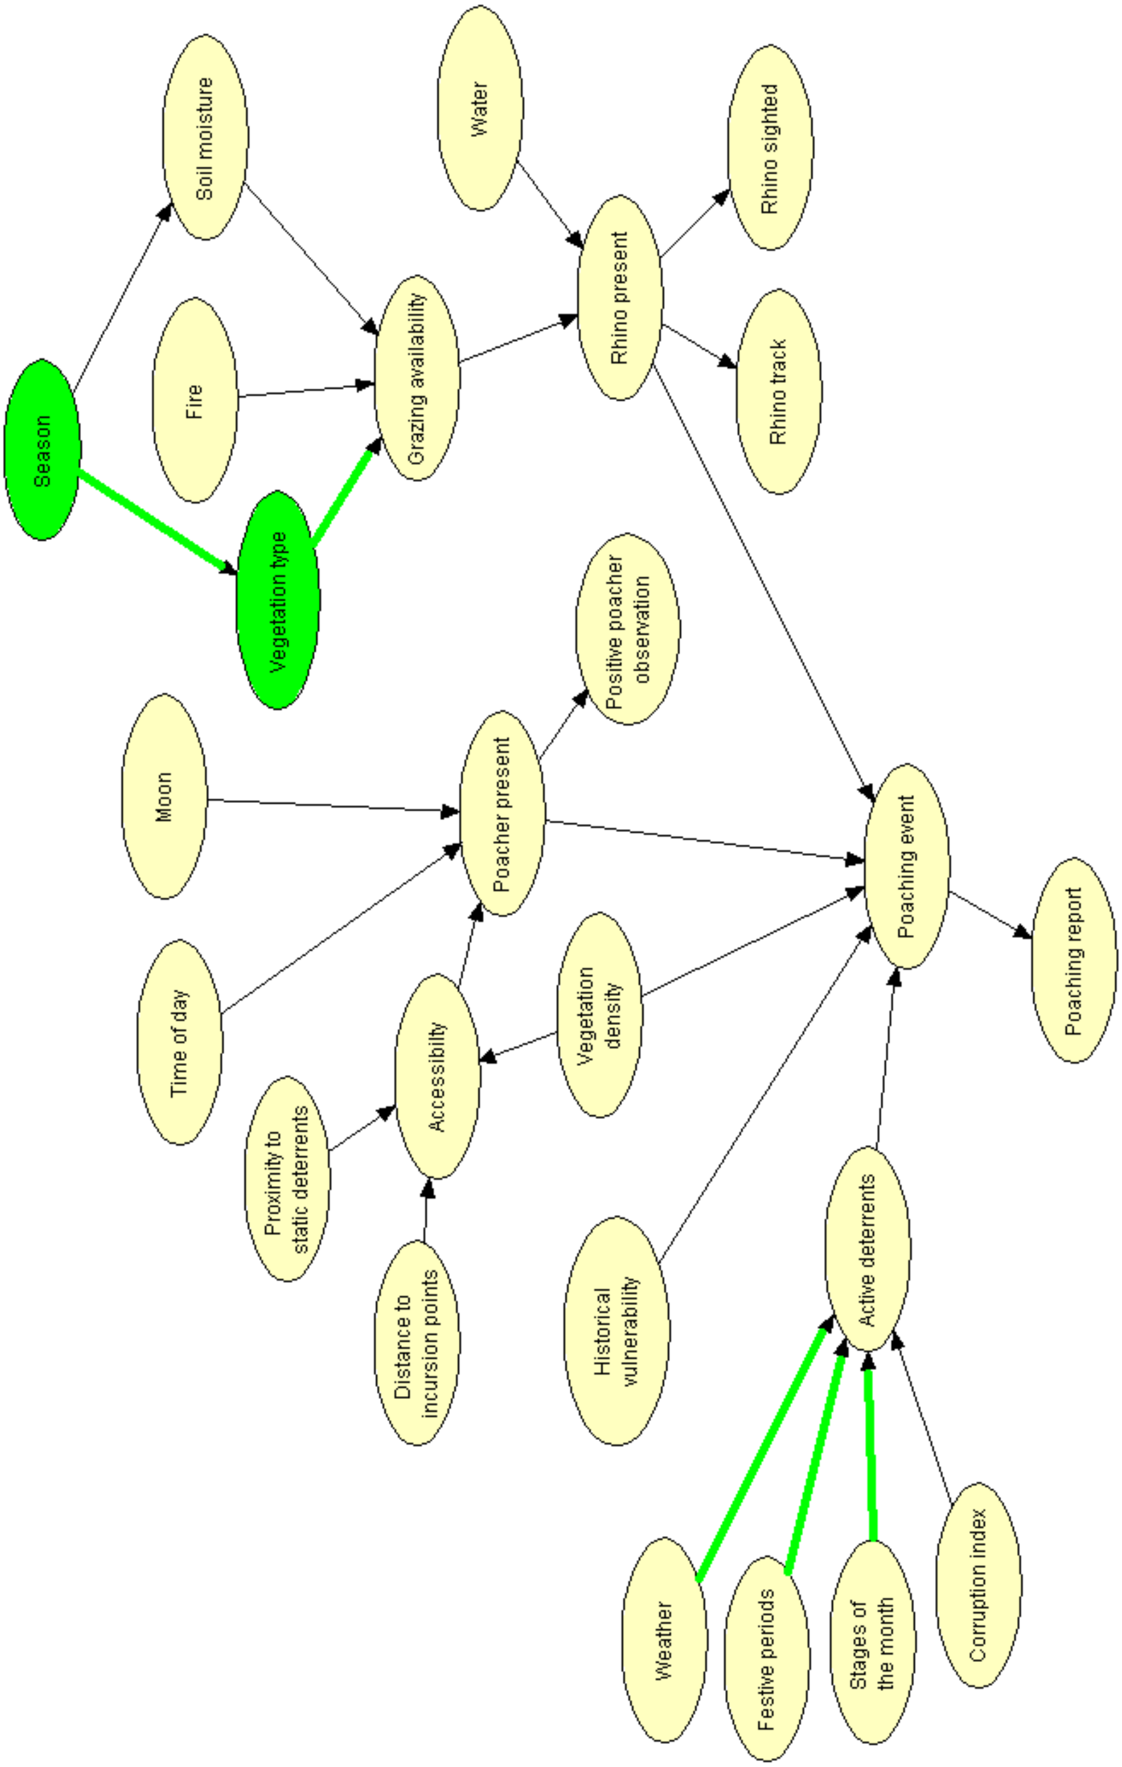

Figure 5.6. Refined model after first round with key experts

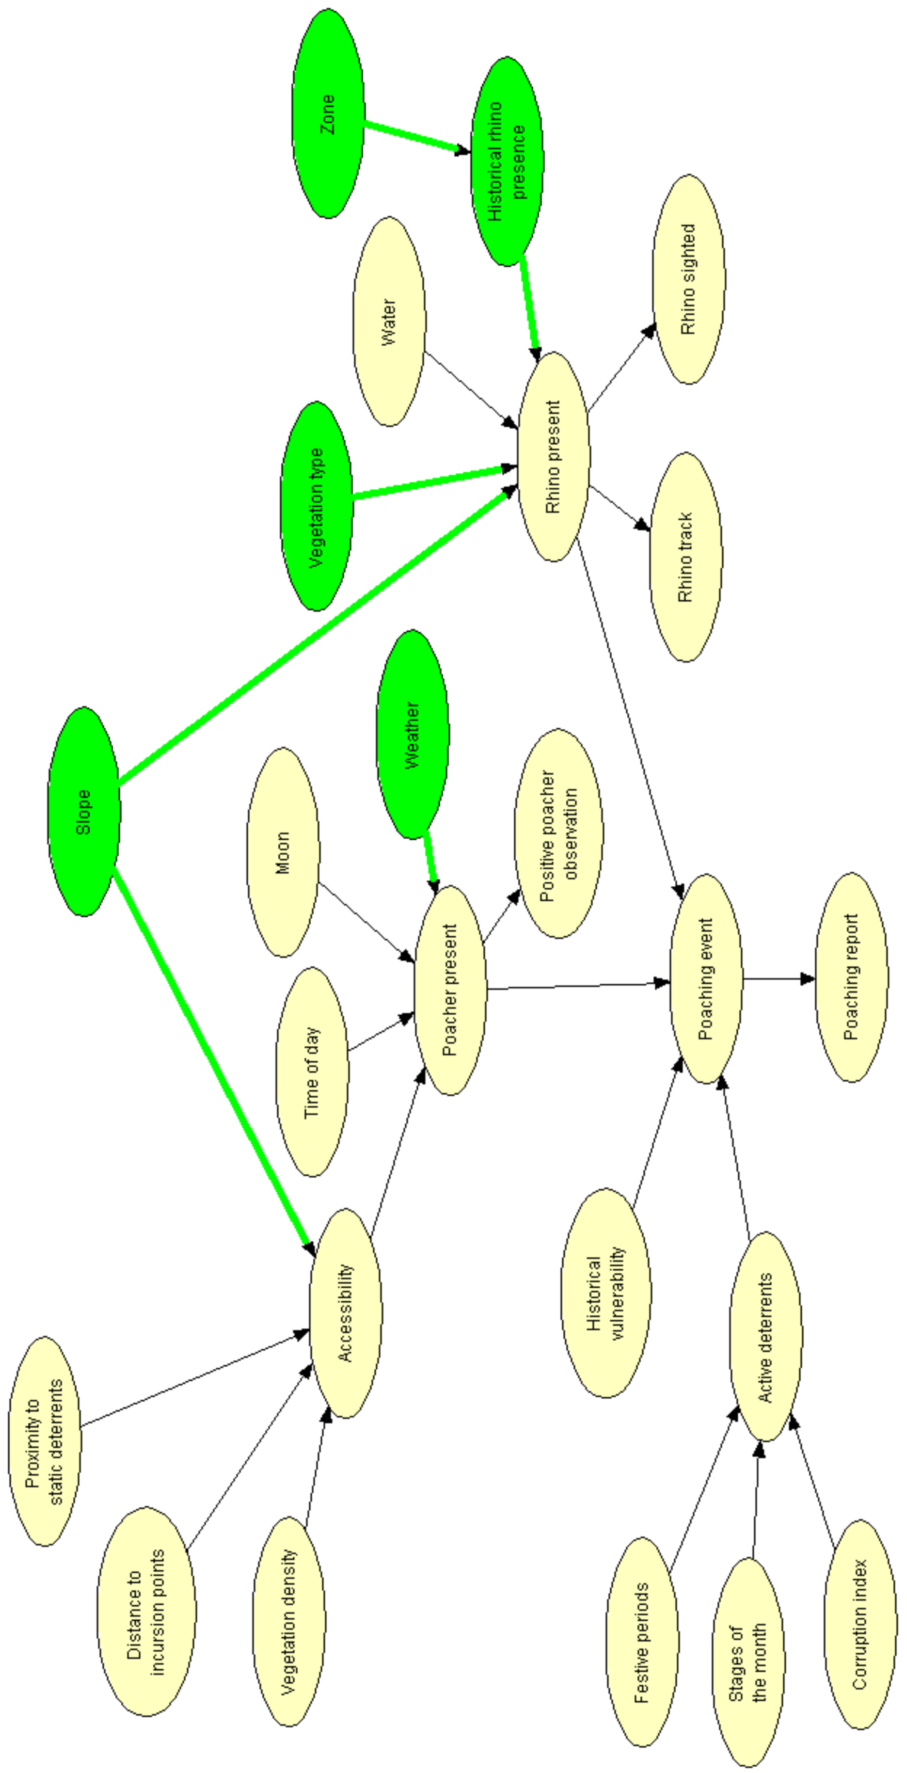

Figure 5.7. Refined model after second round with key experts

The latest network is thus depicted in Figure 5.8. Caution should be taken to not call this the "correct" or the "last" network, as no network will ever be completely correct or finished.

5.6 CONCLUSION

BNs are adept at handling missing or sparse data, as well as with other sources of data including expert input. Real-life problems such as the rhino poaching problem are by nature very sensitive and have a shortage of accurate and complete data, thus an expert knowledge approach is considered. This chapter focusses on the development of a first-iteration BN that was developed by a small group of BN experts and then taken to an expert workshop for evaluation and validation. Developing a BN beforehand saved time and aided the process of obtaining "buy in" from the experts: they could already see a rough model to build on.

A large part of the BN corresponded to the experts' belief and understanding, but some changes were made. Many experts could not attend the workshop due to time and availability constraints, thus one-on-one interviews were set up after the workshop with specific individuals. With every new expert meeting the model changed slightly, but the changes seemed to converge to a unanimous decision. The evolution of the network is shown and discussed at the hand of changing nodes and edges. The latest network is discussed in detail in the following chapter.

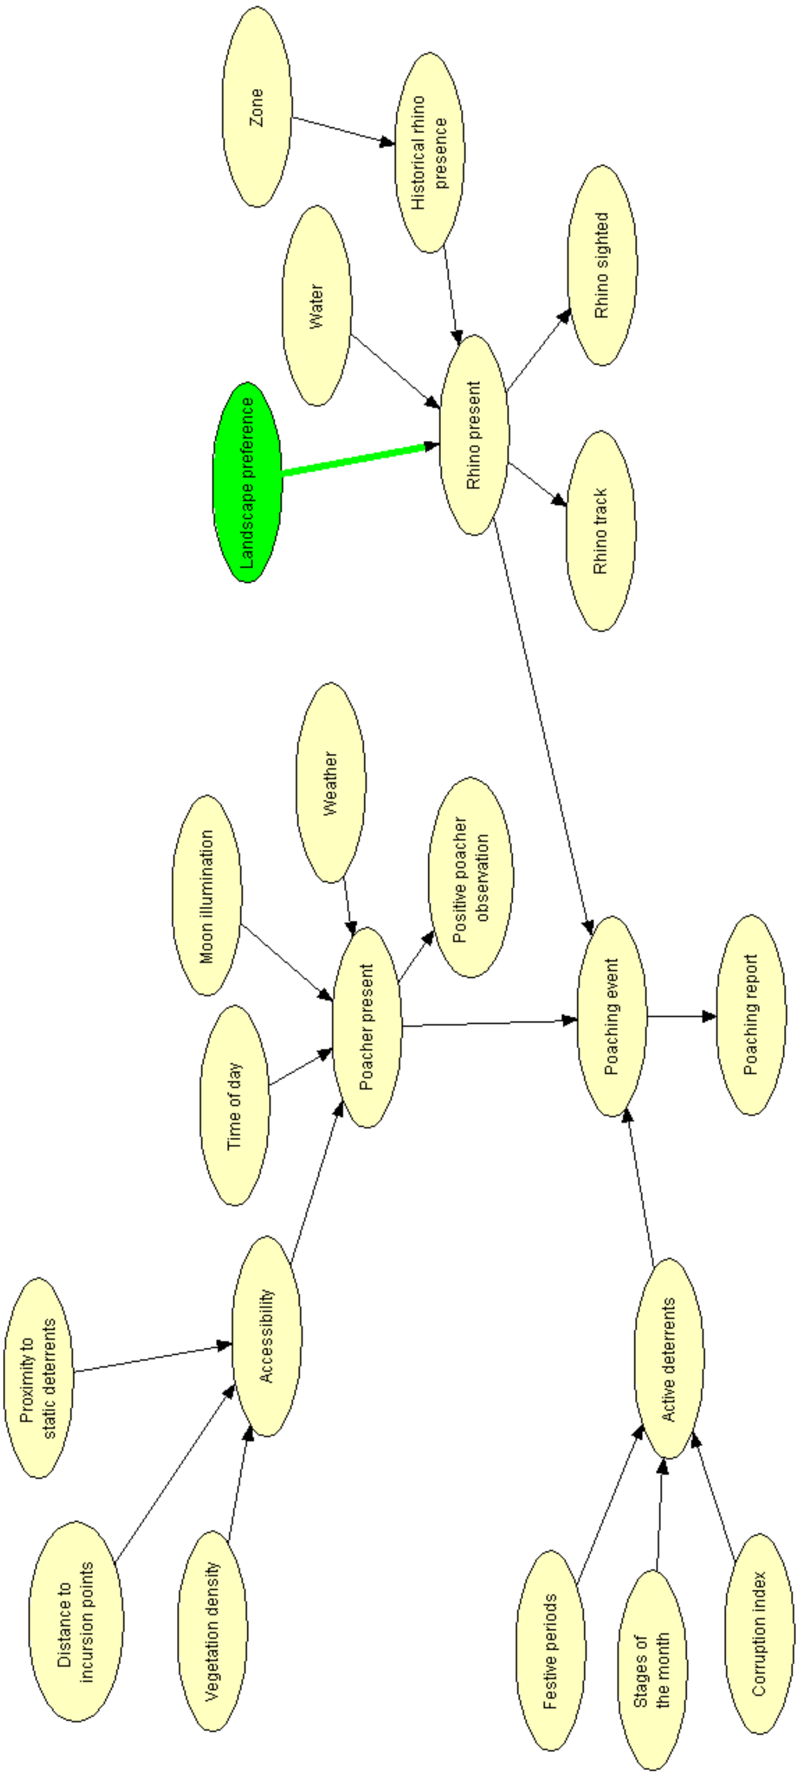

Figure 5.8. Latest iteration of the model

CHAPTER 6 THE FINAL MODEL

6.1 INTRODUCTION

The previous chapter described the process of evaluating an initial causal model with the help of experts to obtain a more refined model. In this chapter the attention is focused on how the different aspects of the model are brought together by the experts' input, as well as the relationships between those aspects. These relationships are causal and are based on the relative conditional probabilities. Causality states that X implies Y, and the experts can say that A causes B and C causes D, and that they believe there is a high likelihood that A influences C.

In this chapter the framework is broken up into parts where it is examined if the probabilities can be found in terms of the distributions, and if the evidence shows that these likelihoods are accurate. The choice of nodes as well as the choice of states is explained, and the conditional probability tables (CPTs) are presented and discussed. The parts are then put back together to obtain a holistic view of the problem. The model is tested against the experience of individuals.

6.1.1 The big picture perspective

Figure 6.1 shows a simple causal model consisting of four nodes, namely *Active_deterrents*, *Poacher_present*, *Rhino_present*, and *Poaching_event*. These are the four main building blocks of the rhino poaching problem as became evident during the course of the study. The arrows point towards the goal or aim of the model, namely *Poaching_event*. The structure of this simple model is the result from consultation with experts. It follows from Routine Activity Theory (RAT) that for a crime to occur (poaching attack - *Poaching_event*), three conditions have to be met [81]. There needs to be (1) a victim or target (the rhino - *Rhino_present* = "True"), (2) a criminal (the poacher - *Poacher_present* = "True"), and (3) the absence of a capable guardian (ranger - *Active_deterrents* = "Not effective").

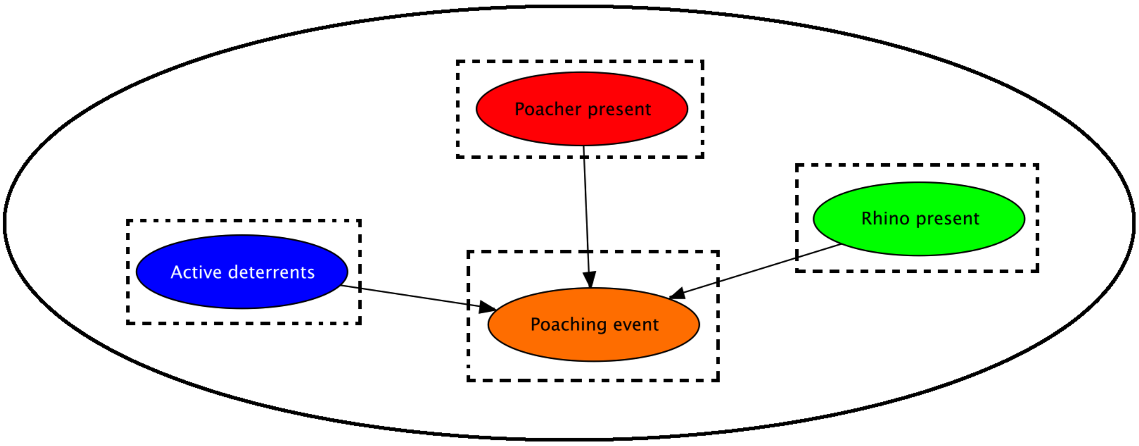

Figure 6.1. The big picture

The simple model represents everything that is known about the rhino poaching problem within the bounds of the Kruger National Park (KNP). The large oval containing the model represents the environment, and everything outside the oval represents everything that is unknown about the rhino poaching problem, as well as everything outside the scope of this study. These unknown elements are of a socio-economical nature, and it is important to understand the problem in its entirety. However, not all aspects are covered in detail in this thesis.

Owing to the initial scarcity and incomplete nature of the data, experts were consulted to form the building blocks as shown in Figure 6.1. The elements in the oval and within the model each represent a subgroup of the complete model after consultation with experts, and is shown in Figure 6.2.

6.1.2 A detail-oriented perspective

Now that the main building blocks of the model have been presented, an in-depth study can be made of the choice of subgroup and each node contained therein. The rest of this chapter will treat the description, motivation, and analysis of each node in the model in Figure 6.2.

The causal model in Figure 6.2 is a Bayesian network (BN) representing the causalities in the model. Causalities are represented by edges and the causalities are associated with conditional probabilities. As is seen in Figure 6.1, the model is partitioned into subgroups, each containing its own small network. The partitioning of the network lends itself naturally towards RAT, and the network will be discussed at the hand of this partition. The words “node” and “variable” are used interchangeably in this work, and represent a random variable in the BN.

Each node in each subgroup will furthermore be discussed according to the following elements:

1. the definition of the node,
2. the motivation for the node,
3. the motivation for the choice of states,
4. the motivation for the choice of conditional probabilities,
5. the motivation for the weightings of the state probabilities conditioned upon the state of the parent nodes,
6. what type of data were used to populate the particular node,
7. how the node was processed.

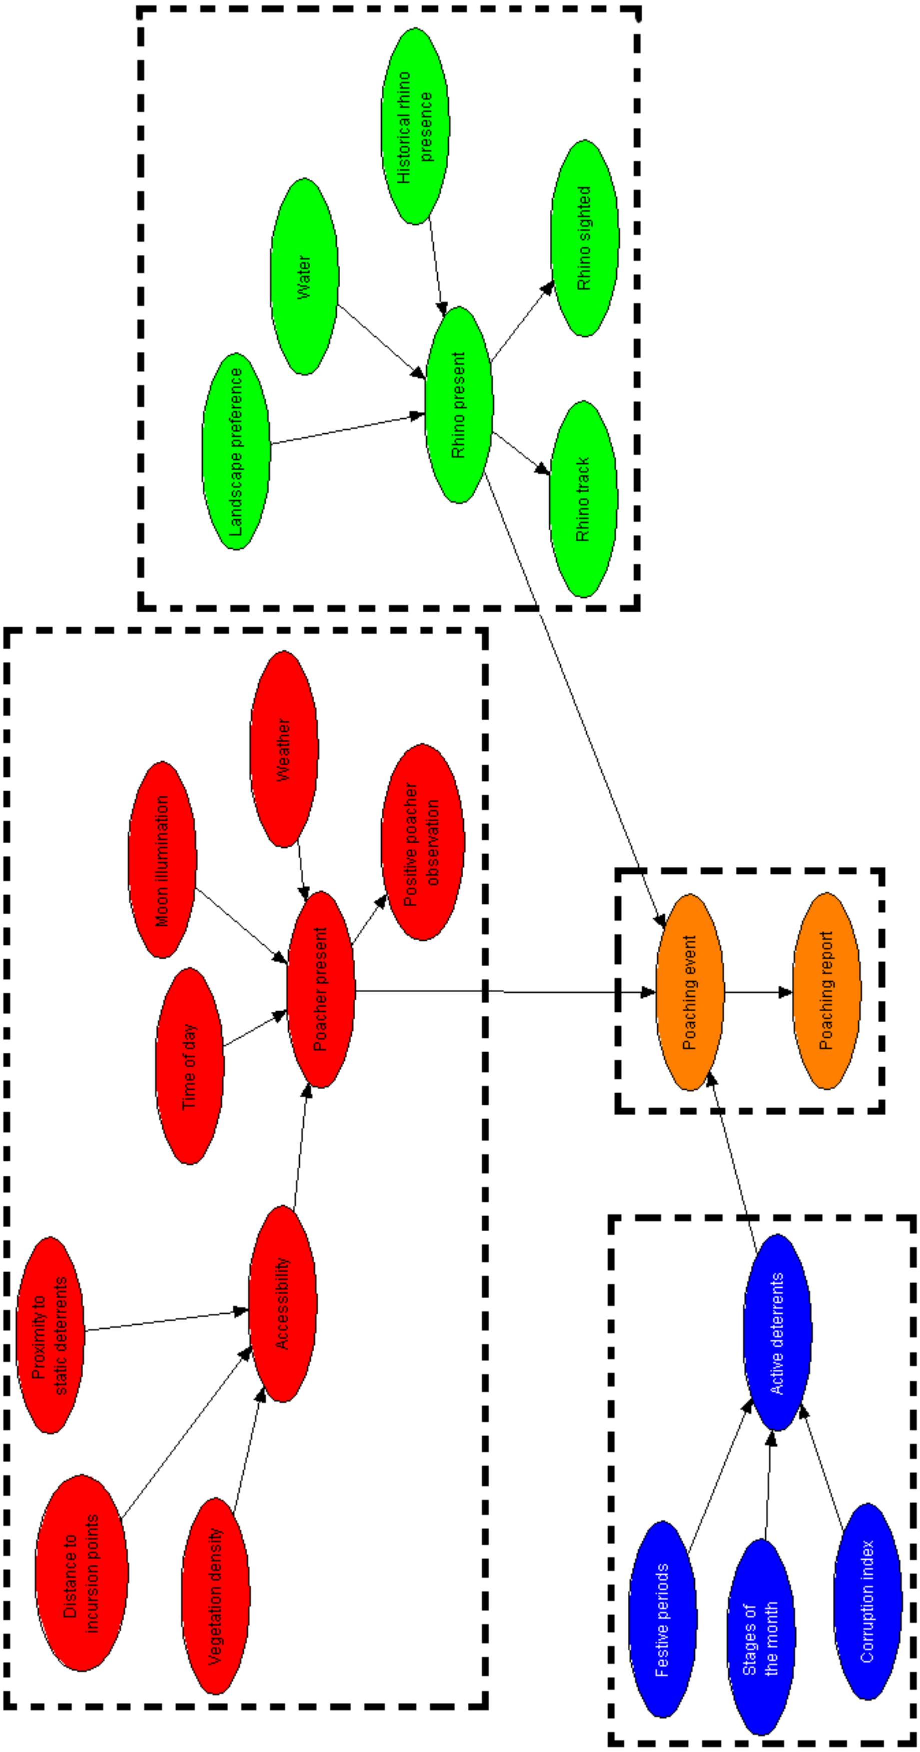

Figure 6.2. Rhino poaching model

6.2 A NOTE ON USER INPUTS

The BN model shown in Figure 6.2 represents snapshots of phenomena that take place within a specific time interval within a specific cell. The variables can be split into time variables and spatial variables. Consider a map of the KNP partitioned into $i = 1, \ldots, M$ cells of size 5×5 kilometres, as agreed upon by the experts, where the ith cell is denoted by C_i as shown in Figure 6.3. Each cell C_i has its own configuration of states (evidence) for the variables of the model. Cell C_i might be close to water, while cell C_{i+1} might be far from water, thus their states for *Distance_to_water* will be different.

The same model is used for each geographically bounded area or cell, only the evidence of the model changes. For prediction of poaching events, a cell size must first be specified. For the dimensions of the KNP it was concluded that 5×5 kilometres would be a satisfactory size to capture enough detail per cell. Each cell is an instance of the BN, thus a cell must be large enough to contain sufficient information to populate the BN. Code was however written to allow for dynamic cell sizes, to allow for future investigations into model performance for different cell sizes.

The user specifies a prediction date and a time of day for which he would like to know the probability of a poaching event, although this will only lighten or darken the map. The user can also decide if he wants to obtain a prediction for a specific location in the KNP (he clicks on a specific area), or if he wants a heat map for the entire KNP. In the event that the user clicked on a specific area, the algorithm locates the specific cell in which the clicked point falls. Otherwise the algorithm starts at the bottom left cell and traverses the entire map, calculating prediction probabilities as it goes along. For the rest of this chapter a single cell will be denoted by C_i.

Obtaining the probability for a single cell is not very useful, as the absolute probability of a cell does not provide us with much information. The relative cell probabilities are much more important, as they tell us what the probability of a cell is in relation to other cells. If all the cells in an area have a very low poaching probability (say, < 0.10), but a cell in that area has a probability of 0.30, then we can say that that cell is much more susceptible to poaching events.

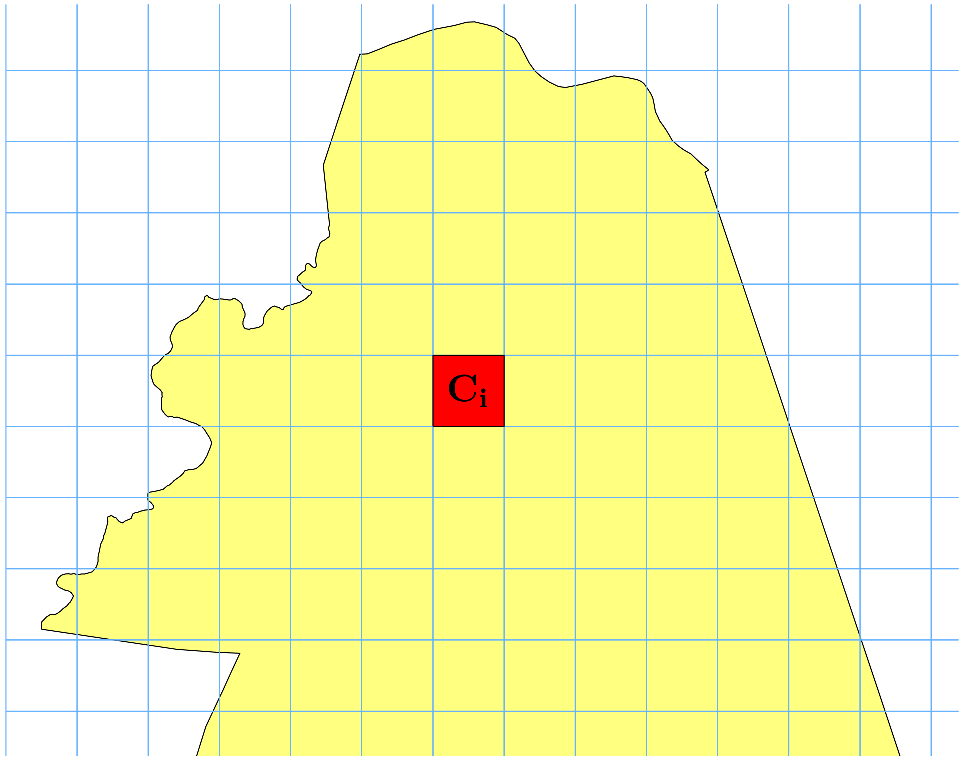

Figure 6.3. Map of KNP divided into 5×5 kilometre cells

6.3 POACHER SUBGROUP

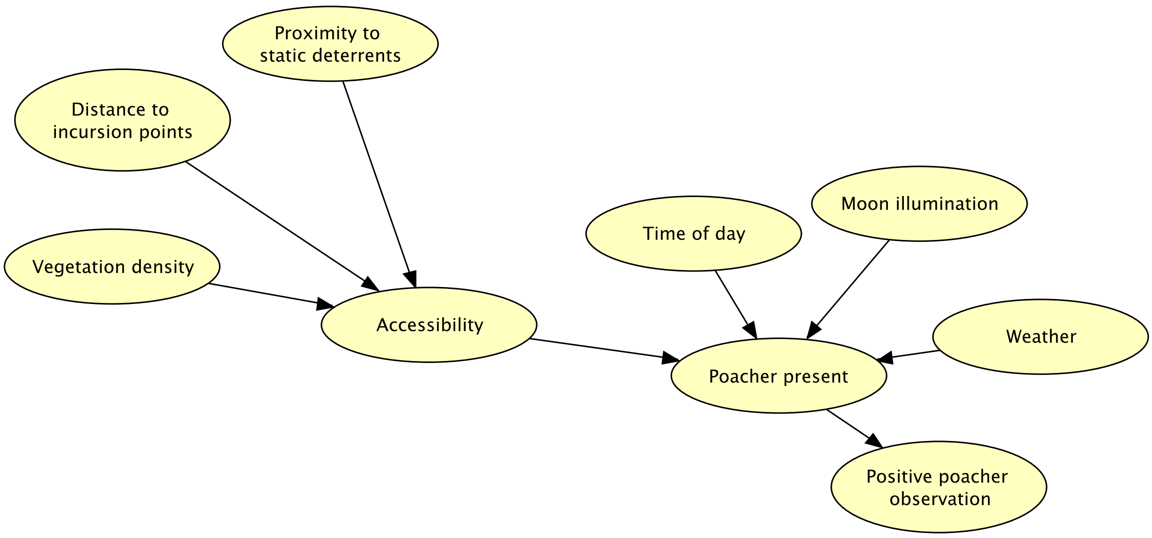

Figure 6.4. Poacher subgroup

The poacher subgroup contains all the available knowledge pertaining to the presence of a poacher. It contains a mixture of spatial as well as temporal variables and is presented in Figure 6.4.

6.3.1 Vegetation density

The goal of *Vegetation_density* is to establish where in the KNP there are areas that are very dense and that could limit the poacher's access to other areas. The states of *Vegetation_density* divide the KNP into areas that are impenetrable and areas that are open, thus the states are "Impenetrable vegetation" and "Open areas". Any states that are not considered as possibly having impenetrable vegetation are classified as "Open areas".

The aim of this variable is to quantify how difficult it is for a poacher to reach a certain area depending on the vegetation. If the vegetation is very dense, it will prevent the poacher from entering or leaving an area. Dense vegetation might also prevent a poacher from entering the KNP at a certain location, or limit the poacher's access to certain areas. However, dense vegetation will camouflage the poacher and make it more difficult to find him. *Vegetation_density* is used as a positive parameter for poachers as it is something that hides their presence from the rangers.

Rhinos prefer open areas, so if an open area is next to an impenetrable area, the poacher might prefer the impenetrable area as that means he can hide and wait for the rhino, undetected. Open areas are detrimental to the poachers as rangers or other individuals can easily see poachers. It is also more difficult for poachers to get away when they are spotted, as there are no trees or bushes to hide in. Impenetrable vegetation will make it easy for the poachers to hide, but it will be extremely difficult for the poachers to get to the rhinos, especially in the dark. White rhinos will also avoid impenetrable vegetation as it is too difficult to enter and exit, and, as they are grazers, there is most likely no suitable vegetation for them there.

This variable does not distinguish between which type of vegetation density is necessarily preferred by a poacher, it simply states the probability of encountering impenetrable vegetation and open areas in the KNP. Table 6.1 illustrates the prior probabilities of this root node. The prior probabilities are only useful if reasoning is performed with the BN where questions are asked about the park as a whole (all the cells). This node will be populated with evidence when the cells are considered individually. Figure 6.5 shows the prior map for the KNP in terms of vegetation density. Red corresponds to impenetrable vegetation and green corresponds to open areas. The prior probability was calculated according to how many cells fall within impenetrable vegetation areas versus how many cells fall within open areas.

A grid size of 5×5 kilometres was used throughout. It is important to obtain prior probabilities to calculate probabilities for all cells even when certain variables are not observable.

At the start of the study no data was available, thus the priors weighted very heavily. In general priors are not that important since we supply evidence to nodes, but it was decided to include it nonetheless for the sake of completeness and reproducibility.

Table 6.1. *Vegetation_density* CPT

States	Probability
Impenetrable vegetation	0.3073
Open areas	0.6927

If a cell contains both impenetrable vegetation and open areas then the largest zone dictates the state of the cell. This instance is shown in Figure 6.6. For this cell the area of the cell covered by impenetrable vegetation is calculated, as well as the total area of the cell covered by open areas. If the impenetrable vegetation area is larger, then that cell is assigned the state of impenetrable vegetation. Figure 6.6 shows both cases: a cell where the impenetrable vegetation has the largest area (1), and a cell where open areas have the largest area (2).

A shapefile that contains land types with different attributes for land type, soil type, altitude, geology, and dominant woody cover was used to determine which cells fall into which state. For *Vegetation_density*, a land types shapefile with different attributes for land type, soil type, altitude, geology, and dominant woody cover was used. It was postulated that there would be very few areas in the KNP with impenetrable vegetation in comparison to open areas. Since there are only two states in *Vegetation_density*, the areas with impenetrable vegetation were identified first, since the open areas would just be what is left on the map.

The dominant woody cover attribute of the shapefile was used to identify the impenetrable vegetation areas. The dominant woody cover is displayed as polygons representing the different areas and types of tree cover. This particular shapefile contains 100 polygons. Woody cover refers to the trees and shrubs in an area, which in turn determines what type of grass will grow in that area. According to Bucini [88], woody cover is defined as the “vertical projection of a woody plant crown”. Woody cover determines many ecosystem functions, such as grazing and browsing availability.

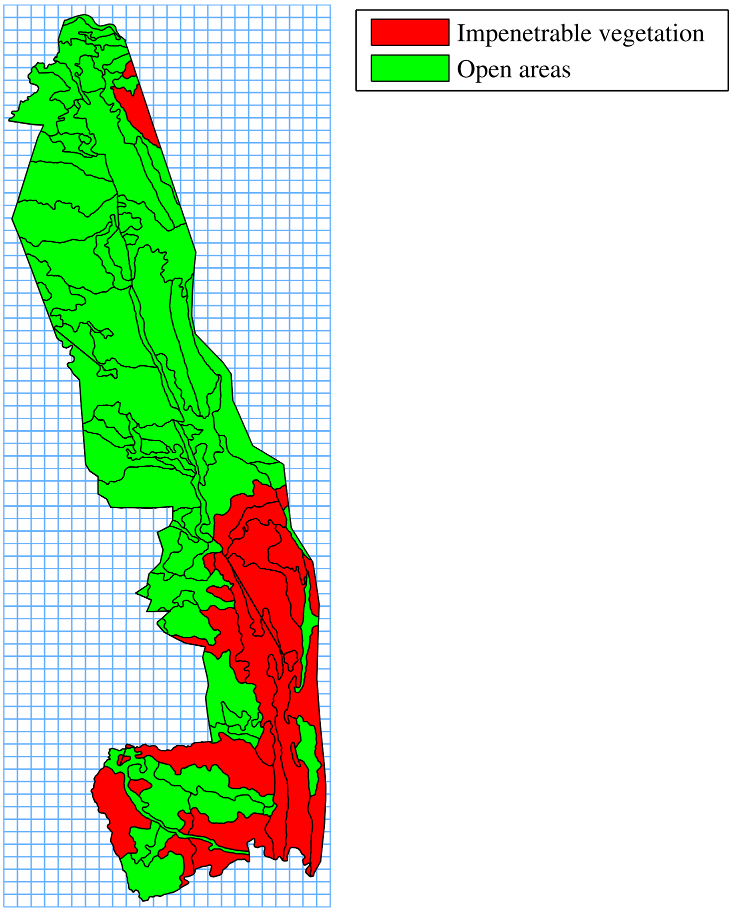

Figure 6.5. *Vegetation_density* prior map

The properties of this attribute are in the form of scientific plant names, thus a lookup table was created in order to identify entries which met the requirements of impenetrable vegetation. Experts and books on woody cover were consulted to ascertain which keywords would constitute impenetrable vegetation. A lookup table was formed with the following keywords: *Dichrostachys cinerea, cinerea, thicket, welwitschii, nigrescens, tortilis*.

The scientific name *Dichrostachys cinerea* refers to the sickle bush (or "sekelbos") which grows in dense thorn thickets [89]. Black rhinos are one of the only species to favour this vegetation type, but for this thesis only white rhinos are considered. Not all of the entries in the land types shapefile are held to the same standard in terms of consistency. Some entries will list "Dichrostachys cinerea" while other will just list "D. cinerea", making it difficult to use these keywords as a lookup table. The single term "cinerea" was thus included to ensure that all cases where it is mentioned are included.

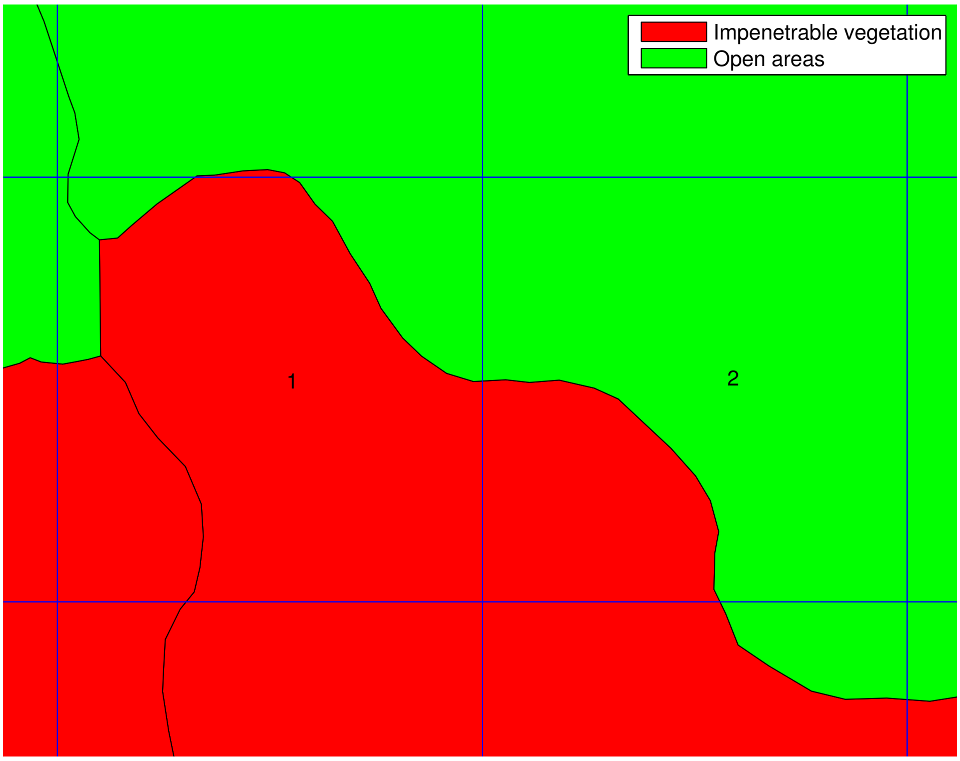

Figure 6.6. *Vegetation_density* prior map

A *thicket* refers to a very dense stand of trees; therefore it was included as a keyword. The terms *welwitschii* and *nigrescens* refer to the *Acacia welwitschii* (Delagoa thorn) and *Acacia nigrescens* (knob thorn or "Knoppiesdoring") respectively [90]. Both types of trees form thickets that can make the area impenetrable. A third type of Acacia, *Acacia tortilis* (umbrella thorn or "haak-en-steek") [90], forms a canopy and can become quite a large tree. Again, the keyword "tortilis" was used to denote both "Acacia tortilis" and "A. tortilis".

The states of each node should undergo the clarity test to verify that the states are mutually exclusive and clear. In order to satisfy the clarity test, the states of a node need to be such that a person with knowledge of past, present, and future events can understand it without additional information [43]. For example, stating that a rainfall node has states "High" and "Low" does not pass the clarity test. A better option would be to state "Higher than the average rainfall" and "Lower than the average rainfall", where the average rainfall is known.

The state "Impenetrable vegetation" refers to vegetation that is so dense that no person can simply walk through it. The state "Open areas" refers to all those areas in the KNP that do not contain impenetrable vegetation. The conditions for the clarity test are thus met.

6.3.2 Distance to incursion points

Incursion points (also called "crossing points") are places where poachers can easily gain entry to the park or exit the park. Incursion points are usually holes in the fence, river crossings, and available pick-up points such as tar roads. This node depicts the distance from the cell to these incursion points. The hypothesis is that the closer a poacher is to an incursion point, the more likely he is to commit a crime, because then he is closer to an exit point as well. The further away from the crossing point, the more likely he is to be detected by rangers.

The states of *Distance_to_incursion_points* divide the KNP into cells that contain an incursion point, cells with a neighbouring cell that contain an incursion point, and cells that are far away from incursion points (none of the cell's neighbours contain an incursion point). There are a number of known crossing points and these points were used to compute the prior probabilities of this root node. Table 6.2 illustrates the prior probabilities. The preference of a poacher towards incursion points will change due to availability and detectability, but for the purpose of this work it was decided to use a static map of the known and most popular incursion points available at the time of the study.

Table 6.2. *Distance_to_incursion_points* CPT

States	Probability
Incursion point in cell	0.0927
Incursion point in neighbouring cell	0.2192
Incursion point far from cell	0.6881

This variable does not rank the probabilities according to a poacher's preference, it simply gives the probability of being a certain distance from an incursion point given a random location in the KNP. The *Distance_to_incursion_points* variable is a positive variable for the poachers, as it influences their willingness to commit a crime.

The probabilities were obtained by using an incursion point shapefile where the incursion points are represented by coordinate pairs in the file. At the time there were only a handful of known incursion points, thus it was hypothesised that most cells would be far from an incursion point. There are two possible ways of approaching this problem: (1) take the midpoints of each cell and calculate the distance to each incursion point, or (2) determine whether the cell or one of its neighbouring cells contain an incursion point. It was decided to use the second option with a cell size of 5×5 kilometres.

All the cells were initially assigned the state of “Incursion point far from cell”, after which further calculations were made. If a cell contained an incursion point, that cell was assigned the state of “Incursion point in cell”. If a cell did not contain an incursion point, its neighbouring cells were calculated. If any of the neighbouring cells contained an incursion point, that cell was assigned the state of “Incursion point in neighbouring cell”.

As expected, there were not many cells that contained an incursion point. The most cells were, as expected, far from an incursion point. The three states pass the clarity test, as the states are clearly defined mutually exclusive intervals on the real line. Owing to the sensitivity of the matter a map of the incursion points can unfortunately not be disclosed.

6.3.3 Proximity to static deterrents

The *Proximity_to_static_deterrents* node quantifies how far a cell is from a static deterrent. In this context, a static deterrent can either be a rest camp, lodge, ranger post, patrol hut, or gate. It is static in that it is at a fixed, although sometimes unknown, location. It is a deterrent because it is believed that poachers prefer to stay away from these locations for fear of being detected. Rhinos may also avoid built-up areas such as camps with considerable human activity, so there may also be no need for the poacher to go there.

This node has a negative impact on the poachers, and also affects the poachers’ accessibility in the park. The proximity to static deterrents is based on the distance that a gunshot can be heard by someone with reasonable skill and expertise. Shot spotters estimate this distance to be between three and four kilometres. The further away from a camp or other structure a ranger is, the more likely he is able to hear a shot. Poachers will not attack close to camps for fear of being captured.

This node denotes the probability of being a certain distance from static deterrents given a random location in the KNP. The distance also depends on the size of the game reserve or park. A small game reserve such as Sabi Sand Game Reserve (also called Sabi Sand Wildtuin (SSW)) will most likely have intervals for "0 – 1 km", "1 - 5 km", and "> 5 km". For a game reserve the size of the KNP, larger intervals have to be chosen. The experts agreed to the same intervals as for the *Distance_to_incursion_points* node (Expert workshop, personal communication, September 29 - October 1, 2014).

Table 6.3 illustrates the prior probabilities for *Proximity_to_static_deterrents* computed from maps of rest camps, rangers posts, patrol huts, and gates. These shapefiles contain elements in the form of point locations, which makes it easy to compute whether or not a cell contains any of the points. The prior probabilities for *Proximity_to_static_deterrents* were calculated in the same way as for *Distance_to_incursion_points*. All the cells were initially assigned the state of "Deterrent far from cell", and then the cells were evaluated in terms of the static deterrent points. If a cell did contain a static deterrent, that cell was assigned the state of "Deterrent in cell". If a cell did not contain a static deterrent point, its neighbouring cells were calculated. If any of the neighbouring cells contained a static deterrent point, then that cell was assigned the state of "Deterrent in neighbouring cell". This variable passes the clarity test, as the states are clearly defined mutually exclusive intervals on the real line. Owing to the sensitivity of the matter a map of the static deterrents can unfortunately not be disclosed.

Table 6.3. *Proximity_to_static_deterrents* CPT

States	Probability
Deterrent in cell	0.1525
Deterrent in neighbouring cell	0.5130
Deterrent far from cell	0.3345

6.3.4 Accessibility

The *Accessibility* node combines the previous three nodes. *Accessibility* quantifies how easy or difficult it is for a poacher to enter the park, or move around in the park, and get to the rhinos. *Accessibility* is influenced by *Vegetation_density*, *Distance_to_incursion_points*, and *Proximity_to_static_deterrents*. This node does not affect whether there is a rhino present or not, but it influences whether or not the poacher can get to the rhinos. The node contains factors that do not change readily.

The *Accessibility* node has two states, namely "Low" and "High". This informs us whether the cell is accessible by poachers or if it is more difficult to access. "Low" means that the chance for a poacher to enter the cell is zero to slim, while "High" means that the poacher's chance of entering the cell is very high. A previous version of the model included a third state, "Medium", but it became ineffective to try and establish clear boundaries between "Low" and "Medium" and between "Medium" and "High".

Table 6.4 illustrates the probabilities for this node given its parent nodes. This specific CPT has a total of $2 \times 3 \times 3 \times 2 = 36$ entries corresponding to the number of states of each node, and therefore only a few entries are shown below. The representation of the probability of *Accessibility* conditioned on its parent variables can be written as:

$$P(Accessibility|Vegetation\_density, Distance\_to\_incursion\_points, Proximity\_to\_static\_deterrents).$$

Table 6.4. *Accessibility* CPT

States				
Vegetation_ density	*Distance_to incursion_points*	*Proximity_to static_deterrents*	*Accessibility*	Probabilities
Impenetrable	In cell	In cell	High	0.9640
Open	In cell	In cell	High	0.9100
Impenetrable	In neighbouring cell	In cell	High	0.9748
Open	In neighbouring cell	In cell	High	0.9370
⋮	⋮	⋮	⋮	⋮
Impenetrable	Far from cell	Far from cell	High	0.1800
Open	Far from cell	Far from cell	High	0.4500

It is important to note that, although four decimal numbers are used throughout the probability elicitation (and the thesis), the experts mostly gave answers in the form of "0.6", "0.2", "half of the time", "a quarter of the time", *etcetera*. The four decimal numbers are important for the CPT calculations where the weightings are multiplied.

The ideal situation for a poacher in terms of accessibility would be when the vegetation density is open (easy access through the vegetation), the distance to incursion points is at a minimum (easy to gain entry to the park or to exit the park), and the proximity to static deterrents is at a maximum (far away from tourists and rangers). Owing to the size of the CPT, it is difficult to elicit probabilities from experts in this form. The CPT was broken up into smaller, more manageable tables.

Table 6.5 illustrates the simplified CPT for the case where $Accessibility$ = "High" and $Vegetation\_density$ = "Open areas". Table 6.6 illustrates the simplified CPT for the case where $Accessibility$ = "High" and $Vegetation\_density$ = "Impenetrable vegetation". $Accessibility$ has two states, one of which is encapsulated in these tables ($Accessibility$ = "High"). The probabilities for the same configuration but for $Accessibility$ = "Low", is the complement of the tables in Table 6.5 and Table 6.6 respectively. These probabilities are written in the $Accessibility$ CPT as shown in Table 6.4. The CPTs of some of the following variables were reviewed by experts to give them the opportunity to adjust the calculated values.

Table 6.5. Simplified CPT for $Accessibility$ = "High", $Vegetation\_density$ = "Open areas"

Proximity_to static_deterrents		*Distance_to_incursion_points*		
		In cell	In neighbouring cell	Far from cell
		1.0000	*0.7000*	*0.5000*
In cell	*0.1000*	0.0900	0.0630	0.0450
In neighbouring cell	*0.5000*	0.4500	0.3150	0.2250
Far from cell	*1.0000*	0.9000	0.6300	0.4500
Vegetation_density.Open areas				

Table 6.6. Simplified CPT for $Accessibility$ = "High", $Vegetation\_density$ = "Impenetrable vegetation"

Proximity_to static_deterrents	*Distance_to_incursion_points*		
	In cell	In neighbouring cell	Far from cell
In cell	0.0360	0.0252	0.0180
In neighbouring cell	0.1800	0.1260	0.0900
Far from cell	0.3600	0.2520	0.1800
Vegetation_density.Impenetrable vegetation			

A weighting between the tables was established so that the expert only had to specify one probability and a few weights. As already mentioned, the accessibility would be the highest if the vegetation is open and if the poacher is close to an incursion point, yet far from a static deterrent. This configuration of states is then chosen as the "baseline probability" in the *Vegetation_density* = "Open areas" CPT. The expert gives an estimate for this probability, in this case the answer was 0.9000. All other calculations are then performed from this baseline. The expert assigns a weighting to each state of each variable present in the CPT. Since *Distance_to_incursion_points* = "In cell" and *Proximity_to_static_deterrents* = "Far from cell" is the best that we can do to obtain a high accessibility, both of these states have a weighting of one. The rest of the weightings are as shown in Table 6.5. The CPT is populated by multiplying the weights of the corresponding states with the baseline. For instance, the entry for *Distance_to_incursion_points* = "In neighbouring cell" and *Proximity_to_static_deterrents* = "In neighbouring cell" would be $0.7000 \times 0.5000 \times 0.9000 = 0.3150$ as is illustrated in Table 6.7. The last weight that is asked of the expert is the relation between the states of *Vegetation_density* when *Accessibility* = "High". The answer in this case is 0.4000, thus the CPT for Table 6.6 is computed by multiplying Table 6.5 with 0.4000.

Specifying only one probability could propagate potential bias through the entire CPT, but sometimes it seemed to work better to rather ask for ranges and use comparative language: Is this more likely to happen than that? If this is true, is this more likely than that? 1 out of 10 times, or 1 out of 5 times? This is quite subjective, but the true refining of the answers only happened once the experts saw the model in action - they could see how their answers yield certain probabilities. The author could ask if the pattern yielded by the model is correct, rather than if the values are correct.

Table 6.7. Example of calculations for *Accessibility* = "High", *Vegetation_density* = "Open areas"

Proximity_to static_deterrents		*Distance_to_incursion_points*		
		In cell	In neighbouring cell	Far from cell
		1.0000	*0.7000*	*0.5000*
In cell	*0.1000*	0.0900	0.0630	0.0450
In neighbouring cell	*0.5000*	0.4500	**0.3150**	0.2250
Far from cell	*1.0000*	0.9000	0.6300	0.4500
Vegetation_density.Open areas				

The CPTs presented in this chapter might look different to readers who are familiar with Hugin®, but the CPTs are presented here as they are used in Kevin Murphy's Bayesian Network Toolbox (BNT) for Matlab® [91]. The child nodes are in topographical order from left to right, with the parent node on the right, and the probability column on the utmost right hand side. This will be explained in the following chapter.

6.3.5 Time of day

Carcasses are normally detected days, weeks, or even months after poaching incidents. Game reserves rely 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 researchers rely on information concerning suspicious activity in the park, or if they are certain that a fresh carcass was found. If a fresh carcass is found the necropsy will be able to give a time frame in hours. Owing to the overwhelming number of carcasses found, there is a significant backlog concerning the necropsies. This backlog also leads to less accurate time estimations since necropsies are not performed immediately.

The motivation for the $Time\_of\_day$ node is its influence towards the visibility of a poacher. The time of day is more influential in Sabi Sand Game Reserve (SSW) than in the KNP due to the fact that the poachers come into SSW, search for the rhinos, hunt them, and leave. In the KNP the poachers might spend a few days in the park and bide their time for the opportune moment to poach a rhino.

Initially, this node was partitioned into six-hour intervals (morning, midday, evening, night), but during the expert workshop it became evident that only three states need to be considered, namely "Midday", "Twilight", and "Night". The reason for this specific breakdown is that poachers prefer to hunt during twilight and then make their getaway during the night. "Twilight" refers to the period just before daybreak as well as the period just before nightfall. "Night" refers to the period between twilight before nightfall and twilight before daybreak. The state, "Midday", was included for completeness and includes all the time periods not covered by the previous two states. These states pass the clarity test as they are mutually exclusive and can easily be understood.

These times of the day have equal probability of occurring, as can be seen in Table 6.8. This variable simply states the probability of being in a certain state at any given time of day and does not distinguish between poacher preferences.

Table 6.8. $Time\_of\_day$ CPT

States	Probability
Midday	0.3333
Twilight	0.3333
Night	0.3333

6.3.6 Moon illumination percentage

From various communications it became evident that the phase of the moon plays an important role in poaching tactics. Initially the variable had four states corresponding to the four major moon phases, but it became clear that the more important aspect is the percentage of the earth that is illuminated, rather than the specific phase of the moon. $Moon\_illumination\_percentage$ is a discrete variable with four states corresponding to four equally wide percentage intervals: $0-25\%, 25-50\%, 50-75\%, 75-100\%$. The phases of the moon regulate these percentages.

$Moon\_illumination\_percentage$ plays a role in how much a person is able to see at night, and is what is referred to when mentioning moon phases playing a role in poaching. Around new moon it is very dark in the park and neither poachers nor rangers can see very well. Poachers will sometimes use this to their advantage to gain access into the park. Full moon is traditionally the time when poachers attack, because of the high illumination and the fact that they can see better.

This node however does not give probabilities for which moon phase is preferred, it simply gives the prior probability of $Moon\_illumination\_percentage$ being in a certain state on a random date. Calculating the moon illumination percentage is a complex calculation for which an open-source software package called *PyEphem* was used. This package contains several functions for astronomical calculations relating to moon phase, date, and time. Inputting a date yields the moon illumination percentage for that date. All dates between 1 January 2000 and 31 December 2015 were used as input to classify the dates according to moon illumination percentage. The ratio was calculated and used as a prior probability for $Moon\_illumination\_percentage$ as presented in Table 6.9.

Table 6.9. *Moon_illumination_percentage* CPT

States	Probability
$0-25\%$	0.3306
$25-50\%$	0.1692
$50-75\%$	0.1692
$75-100\%$	0.3310

The moon illumination can be calculated by using Equation 6.1 where x represents the normalised phase of the moon (between -1 and 1). The function I is the amplitude probability density function of a sine wave [92].

$$I(x) = \frac{1}{\pi}\frac{1}{\sqrt{1-x^2}}, \tag{6.1}$$

where

$$\int_0^1 I(x)dx = 1.$$

The proportion of each state in the lunar cycle must be quantified in order to calculate the correct weights. This can be done by evaluating the implied integral at the matching intervals as shown in Equations 6.2-6.3.

$$\int_0^{0.25} I(x)dx = \int_{0.75}^{1} I(x)dx = \frac{1}{3} \tag{6.2}$$

$$\int_{0.25}^{0.5} I(x)dx = \int_{0.5}^{0.75} I(x)dx = \frac{1}{6}. \tag{6.3}$$

If the sine wave is analysed according to how much time is spent in a moon phase, it can be seen that both the first quarter (0-25%) and the last quarter (75-100%) span the largest time period. The middle half (25-50% and 50-75%) are both shorter time periods.

6.3.7 Weather

Weather plays an important role in the poaching strategy. If it rains, the watering holes are full which could attract rhinos. Water also washes away a poacher's scent so that sniffer dogs cannot follow him. Normally, people will seek cover when it is raining, and poachers use this to their benefit. When the rangers seek cover under trees or next to structures, the poachers will stay away from those areas and rather hunt rhinos, or make their getaway.

Very hot and sunny conditions also inhibit the ability of sniffer dogs to follow a trail, as the scent dissipates after about half an hour in very hot weather. Animals usually seek out shade during the hot days, but will come out into the open if the weather is overcast and cool, which is beneficial for the poachers. Windy conditions are also beneficial for poachers. Rhinos cannot see very well, but they have an excellent sense of smell. If a poacher uses the windy conditions to his benefit, he can come close to the rhino without the rhino even realising it.

The *Weather* variable is a discrete variable with four states corresponding to four chosen weather types, namely “Clear”, “Overcast”, “Rainy”, and “Windy”. These four weather types were chosen to simplify the problem, but also to cover all the main weather types in the park.

There are only two main seasons in the KNP: a wet summer and a dry winter. The dry winter season ranges from April to September and results in sparse vegetation which makes game viewing very easy [93, 94]. It is excellent for game drives, but also good for the poachers. It is low season, except for the school holidays, and the camps and roads are less busy.

The wet summer season ranges from October to March. From November onwards there is high rainfall. These summer rains peak in January and February and end in April. The scenery of the KNP is very green during the summer and game viewing opportunities are still plentiful. There are many newborn animals, but wildlife viewing is not as good as during the dry season. Normally the grass is very tall, which makes game viewing difficult.

Table 6.10 shows the prior probabilities for the *Weather* root node. The probabilities were derived from the literature pertaining to traditional weather patterns in the KNP and in Limpopo, as well as from personal experience, as no accurate historical weather records were available at the time. The dry winter season (April to September) comprises six months, as do the wet season. However, there are more clear sunny days than there are rainy days: when it rains it usually rains in the afternoon or at night, while the rest of the day is sunny and clear. Thus a probability of 0.6000 was assigned to clear conditions, as opposed to the 0.2000 for rainy conditions. A probability of 0.1000 was assigned to both purely overcast as well as purely windy conditions, as these conditions are usually grouped with other weather conditions such as rain or sunshine and rarely occur on their own.

Since prior probabilities are considered, Table 6.10 does not yield the probability of which weather

Table 6.10. *Weather* CPT

States	Probability
Clear	0.6000
Overcast	0.1000
Rainy	0.2000
Windy	0.1000

type is preferred by poachers, it simply gives the prior probability of the weather being in a certain state given a random date and time.

An open-source weather Application Programming Interface (API) was used in making poaching predictions. Entering a date and a time into the API yielded an answer of one of the four weather states and could be used for inference.

6.3.8 Poacher present

The *Poacher_present* variable is the culmination of all the previous variables in this group. All the factors that influence whether or not a poacher will be in a particular cell, such as *Accessibility*, $Time\_of\_day$, $Moon\_illumination\_percentage$, and $Weather$, are included. The reason for including this variable is that a poacher has to be present in order for a poaching attack to occur. The representation of the probability of $Poacher\_present$ conditioned on its parent variables can be written as:

$$P(Poacher\_present|Accessibility, Time\_of\_day, Moon\_illumination\_percentage, Weather).$$

This node has two states, namely “False” and “True”. Through discussions with experts it became clear that it is important to distinguish between South African (SA) poachers and Mozambican poachers, but that does not reflect in the current BN. The states need to be mutually exclusive for a node in a BN, hence it was decided to rather just use a binary valuable with false and true states. The reason for the initial split between SA poachers and Mozambican poachers was because of the difference in their modus operandi, as well as the type of weapon they use.

Poacher_present is a hidden variable as it cannot directly be observed. This node expresses whether there is a poacher present in the cell or not. Table 6.11 illustrates the probabilities for this node given its parent nodes. This specific CPT has a total of $2 \times 3 \times 4 \times 4 \times 2 = 192$ entries corresponding to the number of states of each node, and therefore only a few entries are shown below.

Table 6.11. *Poacher_present* CPT

States					
Accessibility	*Time_of_day*	*Moon_illumination_ percentage*	*Weather*	*Poacher present*	Probabilities
Low	Midday	$0-25\%$	Clear	False	0.9997
High	Midday	$0-25\%$	Clear	False	0.9991
Low	Twilight	$0-25\%$	Clear	False	0.9730
High	Twilight	$0-25\%$	Clear	False	0.9100
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
High	Night	$75-100\%$	Windy	True	0.5760

Owing to the size of the CPT, it is difficult to elicit probabilities from experts in this form. The CPT was broken up into smaller, more manageable tables. A weighting between the tables was established so that the expert only had to populate a single small table. This follows from Section 6.3 where the same process was followed as for *Accessibility*.

Table 6.13 illustrates the simplified CPT for the case where *Poacher_present* = “True” and *Accessibility* = “Low”. Table 6.12 illustrates the simplified CPT for the case where *Poacher_present* = “True” and *Accessibility* = “High”. *Poacher_present* has two states, one of which is encapsulated in these tables. The probabilities for the same configuration but when *Poacher_present* = “False”, is just the complement of the tables in Table 6.13 and Table 6.12. These probabilities are written in the *Poacher_present* CPT as shown in Table 6.11.

The weights were set up to give a larger weighting to rainy conditions, a slightly smaller weighting to overcast and windy conditions, and a very small weighting to clear conditions. For each weather condition, the moon illumination percentage was weighted more heavily towards the 75% mark than for the 0% mark.

Table 6.12. Simplified CPT for *Poacher_present* = "True", *Accessibility* = "High"

	Moon_illumination		*Time_of_day* Midday	Twilight	Night
			0.0100	*1.0000*	*0.8000*
Clear	0 – 25%	*0.1000*	0.0009	0.0900	0.0720
	25 – 50%	*0.2000*	0.0018	0.1800	0.1440
	50 – 75%	*0.3000*	0.0027	0.2700	0.2160
	75 – 100%	*0.4000*	0.0036	0.3600	0.2880
Overcast	0 – 25%	*0.2000*	0.0018	0.1800	0.1440
	25 – 50%	*0.4000*	0.0036	0.3600	0.2880
	50 – 75%	*0.6000*	0.0054	0.5400	0.4320
	75 – 100%	*0.8000*	0.0072	0.7200	0.5760
Rainy	0 – 25%	*0.3000*	0.0027	0.2700	0.2160
	25 – 50%	*0.5000*	0.0045	0.4500	0.3600
	50 – 75%	*0.7000*	0.0063	0.6300	0.5040
	75 – 100%	*1.0000*	0.0090	**0.9000**	0.7200
Windy	0 – 25%	*0.2000*	0.0018	0.1800	0.1440
	25 – 50%	*0.4000*	0.0036	0.3600	0.2880
	50 – 75%	*0.6000*	0.0054	0.5400	0.4320
	75 – 100%	*0.8000*	0.0072	0.7200	0.5760
Accessibility.High					

6.3.9 Positive poacher observation

The *Positive_poacher_observation* node has states "False" and "True", depending on reports of positive poacher sightings by either game rangers or tourists, or from tracks that were observed. The node gets its information from the *Poacher_present* node as the presence of a poacher in the vicinity increases the likelihood of observing a poacher in a cell. Possible observations include alarms, camera

Table 6.13. Simplified CPT for *Poacher_present* = "True", *Accessibility* = "Low"

		Time_of_day		
	Moon_illumination	Midday	Twilight	Night
Clear	0 – 25%	0.0003	0.0270	0.0216
	25 – 50%	0.0005	0.0540	0.0432
	50 – 75%	0.0008	0.0810	0.0648
	75 – 100%	0.0011	0.1080	0.0864
Overcast	0 – 25%	0.0005	0.0540	0.0432
	25 – 50%	0.0011	0.1080	0.0864
	50 – 75%	0.0016	0.1620	0.1296
	75 – 100%	0.0022	0.2160	0.1728
Rainy	0 – 25%	0.0008	0.0810	0.0648
	25 – 50%	0.0014	0.1350	0.1080
	50 – 75%	0.0019	0.1890	0.1512
	75 – 100%	0.0027	0.2700	0.2160
Windy	0 – 25%	0.0005	0.0540	0.0432
	25 – 50%	0.0011	0.1080	0.0864
	50 – 75%	0.0016	0.1620	0.1296
	75 – 100%	0.0022	0.2160	0.1728
Accessibility.Low				

traps, fence alarms, and spoor. There are very little data for sightings, but there are some data for tracks, hence all these possible factors were included into one node. Table 6.14 shows the probability table.

If a poacher is not present, a ranger might mistake a tourist for a poacher (*Poacher_present* = "False", *Positive_poacher_observation* = "True"), thus the corresponding probability is slightly greater than zero. If a poacher is present, he will hide and try not to be observed (*Poacher_present* = "True", *Positive_poacher_observation* = "True"), thus the corresponding probability for observed is lower

Table 6.14. *Positive_poacher_observation* CPT

States		
Poacher_present	*Positive_poacher_observation*	Probabilities
False	False	0.9500
True	False	0.6000
False	True	0.0500
True	True	0.4000

than for not observed.

6.4 RHINO SUBGROUP

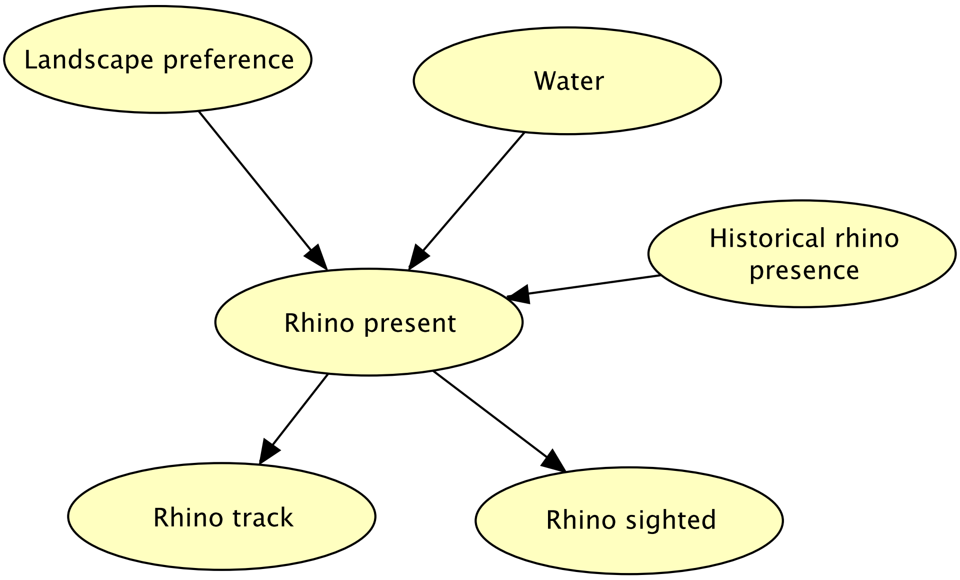

Figure 6.7. Rhino subgroup

The rhino subgroup contains the available knowledge pertaining to the presence of a rhino in a cell. The subgroup's structure is presented in Figure 6.7.

6.4.1 Landscape preference

Rhinos need food to survive, so the rationale is that where the food is will dictate where the rhinos are. 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 quite selective in their grazing habits and they prefer grass types that are more palatable that tends to grow in shade [95].

White rhinos are grazers, thus they only eat grass, unlike black rhinos who are browsers and eat the leaves of low trees and shrubs (and in some cases, thorny plants). The dominant woody cover gives an indication of the types of trees and woody cover in an area, which in turn indicates which grass types are present [95]. Performing an in-depth study into the vegetation preferences of white rhinos proved to be a daunting task, and one that was considered to be beyond the scope of this project.

Communications with experts in wildlife management and ecology revealed that rhinos have an affinity towards muddy pans and waterholes where they like to wallow, especially on hot days. They also like to rest near drainage lines when it is cold. Adding all of these variables into the rhino subgroup would make this subgroup a study on its own. After consulting various sources, amongst others the works by Pienaar *et al.* [95–97], it was decided to combine all of these variables into a single variable called $Landscape\_preference$. The goal of the variable is to inform us where the rhinos are. It does not matter which type of grass they prefer, or which types they avoid, but it matters in which cells they will most likely be present. $Landscape\_preference$ thus encapsulates the vegetation type, the grazing availability, as well as any other habitat preference the rhinos might have. Adding the influence of natural occurrences such as drought, fire, and rain could improve the predictive performance of the model and was considered in a previous version of the model. However, it was decided to remove these variables as inferring the relationships between them proved to be a challenging task.

The states for this variable are “Avoid”, “Neutral”, and “Prefer”. The preferred areas are those areas that contain the best grazing, muddy pans to wallow in, or shade to rest in. The avoided areas are areas where they will absolutely not go for lack of water, the presence of dense vegetation, or rocky outcrops. The neutral areas are the areas that are not preferred, but also not avoided. If over-grazing occurs in the preferred areas or if there was a fire and fresh new leafs form, the rhinos will be attracted to the neutral areas.

The papers of Pienaar *et al.* [95–97] were used as the main reference together with a land types shapefile containing information about the woody cover, the geology of areas, as well as the soil types. It is important to note that the papers of Pienaar *et al.* were written between 22 and 24 years before this thesis, and it is interesting that some of the preferred landscapes and some of the avoided landscapes have remained the same after the 24 years.

Pienaar *et al.* wrote two papers about the landscape preferences of white rhinos in the KNP. The first paper contains information about the landscape preferences in the southern part of the KNP [96] and the second paper about the landscape preferences in the central and northern parts of the KNP [97]. The research of both papers was conducted using yearly aerial census data of white rhino presence together with a chi-square test to ascertain landscape preference or avoidance.

The southern region of the KNP is fenced in by the Sabie River in the north. The central and northern regions are thus the area north of the Sabie River. According to Pienaar [96], the KNP is zoned into 35 landscapes. Of these 35 landscapes, one was identified in the south as the preferred landscape of white rhinos, and two landscapes that are avoided by white rhinos. The preferred landscape in the south is "moderately undulating granitoid plains with *Combretum zeyheri* woodland" [96]. The two landscapes that are avoided are "Low granitoid mountains with *Combretum apiculatum* bushveld" and "Granitoid lowlands with *Acacia grandicornuta* tree-savanna" [96].

The 35 landscapes were again used in Pienaar [97] to determine the landscape preference of white rhinos in the central and northern regions of the KNP. The two preferred landscapes were established to be "Granitoid plains with *Colophospermum mopane* bush savanna" and "Karoo sediment plains with *Acacia welwitschii* tree savanna" [97]. The six avoided landscapes were found to be "Basaltic plains with *Colophospermum mopane* shrub savanna", "Gabbroic plains with *Colophospermum mopane* shrub savanna", "Calcitic plains with *Colophospermum mopane* shrub savanna", "Alluvial plains with *Acacia albida* bush savanna", "Sandy plains with *Baphia massaiensis* bush savanna", and "Andesitic plains with *Combretum collinum* shrub savanna" [97]. Table 6.15 contains the common names for these tree types.

These keywords and phrases for avoided and preferred tree types were matched to keywords and phrases contained in the available land types shapefile. The age of the shapefile can be traced to 2007, yielding an age difference between the shapefile and the papers by Pienaar of 13 to 15 years. However, the

Table 6.15. Common tree names

Scientific name	Common name
Acacia albida/Faidherbia albida	Ana tree (“Anaboom”)
Acacia grandicornuta	Horned thorn (“horingdoring”)
Acacia welwitschii	Delagoa thorn (“Delagoa doring”)
Baphia massaiensis	Sand camwood (“Sandkamhout”)
Colophospermum mopane	Mopane or turpentine tree (“mopane”)
Combretum apiculatum	Red bush willow (“rooiboswilg”)
Combretum collinum	Bushwillow (“boswilg”)
Combretum zeyheri	Large-fruited bush willow (“raasblaar”)

correspondence between the older maps and the current conditions should still be significant enough for the purpose of the model. Shapefiles of the terrain and vegetation are not captured regularly, and a 2007 shapefile is currently the latest version. As already mentioned, the landscapes do shift over time, hence the maps will not look exactly the same after 13-15 years. Table 6.16 illustrates the prior probabilities for $Landscape\_preference$ and Figure 6.8 illustrates the prior map for $Landscape\_preference$ based on the available information. The reader is referred to the papers by Pienaar [96] and [97] to compare the maps.

Table 6.16. $Landscape\_preference$ CPT

States	Probability
Avoid	0.2610
Neutral	0.4644
Prefer	0.2746

6.4.2 Water

Water is vital to the survivability of rhinos. White rhinos will drink water daily to twice daily in wet seasons and undertake long journeys in dry seasons to get to water [95]. A cell with a water source will have a higher probability of having a rhino present than a cell without a water source. The distance to a water source is more influential in the KNP because the KNP does not have as much water as, say

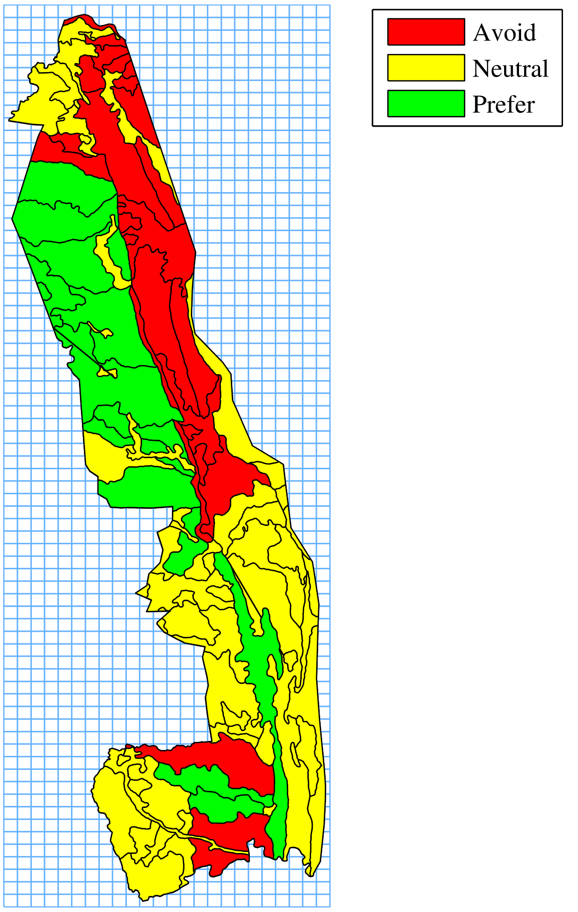

Figure 6.8. Preferred landscapes in the KNP

SSW. Animals will be stationary and close to water during the day, and at night they will lay down in the drainage lines.

The calculations for *Water* were performed by using a water point shapefile, as well as a shapefile containing all the main rivers flowing through the KNP. The water point shapefile was created by merging the shapefiles of waterholes, springs, water points, water pans, and water dams. These are all point coordinates, thus it was straightforward to merge it into a single shapefile. It is assumed that all of the sources (except for boreholes) are suitable for consumption by rhinos. The seasonal effect of water, which can greatly affect water availability, is ignored here. It was decided to use an "average" water availability variable as obtaining data pertaining to water availability during different stages of drought is a study on its own.

The states of *Water* were originally called “Far from water” and “Close to water”, where “Far” and “Close” denoted whether or not a cell was within five kilometres from a water source. Throughout this project, 5×5 kilometre cells are considered. If a cell contains a water source that is on a corner of that cell, it can be argued that the maximum distance from that water source to a point on the diagonally opposite corner is more than five kilometres ($\sqrt{5^2+5^2} = 7.0711$ kilometres), thereby rendering the state’s definition invalid. It was thus decided to call the states ‘Water source not in cell” and “Water source in cell”. This node passes the clarity test, as the states are mutually exclusive and do not overlap.

The closer a cell is to a water source, the higher the probability of a rhino being present, and consequently the higher the probability of a poaching attack taking place. Each cell in the grid is checked if it contains a water source. The water points were checked first, and if there were no water points in a cell, only then were the rivers checked, as the geometry of rivers is captured as lines. All cells are originally marked as “Water source not in cell”. If a water point lies in a cell, the cell is marked as “Water source in cell”, else it stays marked “Water source not in cell”. If all the water points have been checked and the cell is still marked as “Water source not in cell”, the rivers are checked to see if any one of them crosses the cell. If a match is found, the cell is marked as “Water source in cell”, and the procedure for that cell ends. Table 6.17 illustrates the prior probabilities for the *Water* node and Figure 6.9 shows the prior map.

Table 6.17. *Water* CPT

States	Probability
Water source not in cell	0.2215
Water source in cell	0.7785

6.4.3 Zone

The KNP can roughly be divided into three zones, namely the Composite Protection Zone (CPZ), the Joint Protection Zone (JPZ), and the Intensive Protection Zone (IPZ). Each zone has its own difficulties and needs concerning the protection of animals and resources. The zones also have different rhino population densities, adding to the operational difficulties in the various zones.

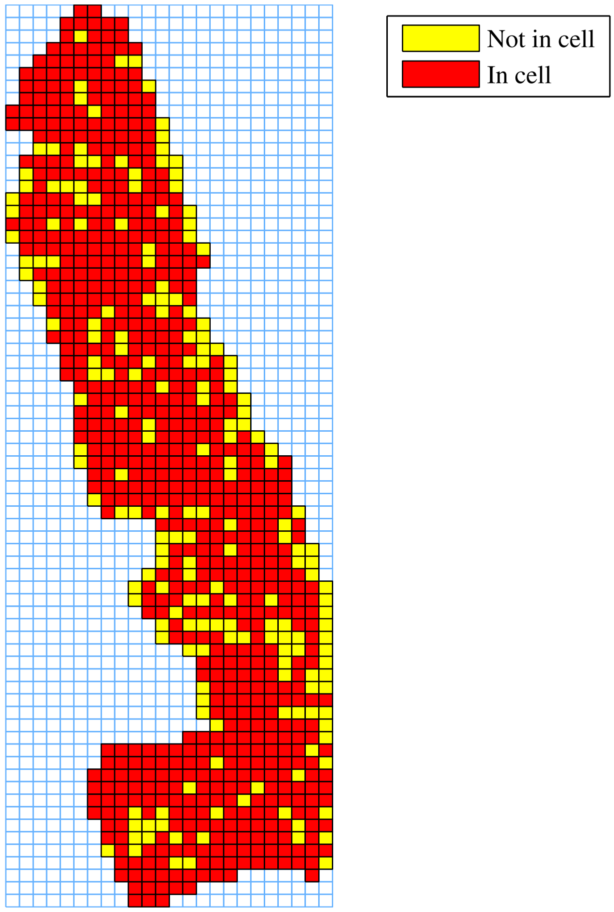

Figure 6.9. *Water* prior map

A previous version of the model contained a variable named *Zone* with states corresponding to the different protection zones. It was used as a switch for *Historical_rhino_presence* to trigger the different presence situations in each protection zone. The variable was however removed as it became clear that the zone is already implied in the *Historical_rhino_presence* data; if the *Zone* variable is used it will be counting the data twice.

6.4.4 Historical rhino presence

The *Historical_rhino_presence* node informs us what the distribution of rhinos in the KNP looked like in the past. If the map for the past four years shows rhinos frequenting a certain area, then we can assume that they will frequent that same area this year as well, as rhinos tend to prefer staying in the same areas. Rhino sighting data from 2011 to 2014 were used to populate the *Historical_rhino_presence* node. The prior probabilities for this node were originally calculated by computing the number of rhino sightings per cell according to the zone it occurred in. After the *Zone* variable was removed, the prior probabilities were calculated according to an estimated number of rhinos per cell.

It is nearly impossible to compute the exact states to use in a node such as this. A decision has to be made over how many states there should be, and also how the states are to be binned. As a start it was decided to use four states, as more states would blow up the CPT to an almost unmanageable size. It was also decided to use $[0,1),[1,3),[3,5),[5,\infty)$ as this would give a good balance between cells where there occurred zero sightings, $[0,1)$, and cells where there occurred an extremely high number of sightings, $[5,\infty)$. The sighting data were used to count how many sightings occurred per cell, and the cells were then categorised according to this number. The resulting prior probabilities can be seen in Table 6.18.

Table 6.18. *Historical_rhino_presence* CPT

States	Probability
$[0,1)$ sightings	0.3921
$[1,3)$ sightings	0.0927
$[3,5)$ sightings	0.0746
$[5,\infty)$ sightings	0.4407

6.4.5 Rhino present

The presence of a rhino plays a cardinal role in the occurrence of a poaching event, and the probability of the presence of a poacher. *Rhino_present* is a culmination of all the factors that influence the possible presence of a rhino. These factors include *Landscape_preference* (Is the habitat suitable?), *Water* (Is water available for consumption?), and *Historical_rhino_presence* (Do rhinos traditionally roam these areas?). *Rhino_present* is a hidden variable as the presence of rhinos cannot always be directly observed.

Rhino_present is a binary variable that takes on the value "True" if a rhino is present in a cell and "False" otherwise. This specific CPT has a total of $3\times2\times4\times2 = 48$ entries corresponding to the number of states of each node. Table 6.19 illustrates the probabilities for this node given its parent nodes. Only a few entries are shown below due to the size of the table. The representation of the probability of

Rhino_present conditioned on *Landscape_preference*, *Water*, and *Historical_rhino_presence* can be written as:

$$P(Rhino\_present|Landscape\_preference, Water, Historical\_rhino\_presence).$$

Table 6.19. *Rhino_present* node CPT

States				
Landscape preference	*Water*	*Historical rhino_presence*	*Rhino present*	Probabilities
Avoid	Not in cell	$[0,1)$ sightings	False	0.9986
Neutral	Not in cell	$[0,1)$ sightings	False	0.9900
Prefer	Not in cell	$[0,1)$ sightings	False	0.9857
Avoid	In cell	$[0,1)$ sightings	False	0.9952
⋮	⋮	⋮	⋮	⋮
Avoid	In cell	$[5,\infty)$ sightings	True	0.0950
Neutral	In cell	$[5,\infty)$ sightings	True	0.6650
Prefer	In cell	$[5,\infty)$ sightings	True	0.9500

The ideal situation for a rhino to be present would be when the landscape is preferred, the rhino is close to water, and there have been many sightings of rhinos in that area over the years. Owing to the size of this CPT, it was broken up into smaller tables such as was done with the CPT of *Accessibility*.

Table 6.20 shows the simplified CPT for the case where *Rhino_present* = “True” and *Landscape_preference* = “Prefer”. Table 6.21 illustrates the CPT for the case where *Rhino_present* = “True” and *Landscape_preference* = “Neutral”, and Table 6.22 illustrates the CPT for the case where *Rhino_present* = “True” and *Landscape_preference* = “Avoid”. *Rhino_present* has two states, one of which is encapsulated in each of these tables. The probabilities for the same configuration but for *Rhino_present* = “False” is the complement of the tables in Table 6.20, Table 6.21, and Table 6.22 respectively. When all three tables, as well as their complements, are combined into one table, the result is the CPT in Table 6.19.

Table 6.20. Simplified CPT for *Rhino_present* = "True", *Landscape_preference* = "Prefer"

Historical rhino_presence		*Water* Not in cell	*Water* In cell
		0.3000	*1.0000*
[0,1) sightings	*0.0500*	0.0143	0.0475
[1,3) sightings	*0.2500*	0.0713	0.2375
[3,5) sightings	*0.7000*	0.1995	0.6650
[5,∞) sightings	*1.0000*	0.2850	**0.9500**
Landscape_preference.Prefer			

Table 6.21. Simplified CPT for *Rhino_present* = "True", *Landscape_preference* = "Neutral"

Historical rhino_presence	*Water* Not in cell	*Water* In cell
[0,1) sightings	0.0100	0.0333
[1,3) sightings	0.0499	0.1663
[3,5) sightings	0.1397	0.4655
[5,∞) sightings	0.1995	0.6650
Landscape_preference.Neutral		

Table 6.22. Simplified CPT for *Rhino_present* = "True", *Landscape_preference* = "Avoid"

Historical rhino_presence	*Water* Not in cell	*Water* In cell
[0,1) sightings	0.0014	0.0048
[1,3) sightings	0.0071	0.0238
[3,5) sightings	0.0200	0.0665
[5,∞) sightings	0.0285	0.0950
Landscape_preference.Avoid		

6.4.6 Rhino track

Rhino_track is an observed variable that states whether a rhino tracking system indicates that tracked rhinos are in a specific cell [25]. *Rhino_track* is influenced by *Rhino_present*: the likelihood of a rhino being present in a specific cell influences whether or not there is a rhino present. If a rhino is not present, the likelihood of a tracking system detecting the presence of a rhino is close to zero. The states are "True" and "False", where "True" represents the state if the presence of tracked rhinos are indicated, and "False" otherwise. Table 6.23 shows the probabilities for these scenarios.

Table 6.23. *Rhino_track* CPT

States		
Rhino_present	*Rhino_track*	Probabilities
False	False	0.9990
True	False	0.2000
False	True	0.0010
True	True	0.8000

The probabilities for the cases where both *Rhino_present* and *Rhino_track* are in the same state are higher than for the other cases. There will be no tracking data if there is no rhino present (P(*Rhino_track* = "False" | *Rhino_present* = "False") = 0.9990), except if the tracking data tracks a different type of animal but concludes that it is a rhino (P(*Rhino_track* = "True" | *Rhino_present* = "False") = 0.0010). The likelihood of tracking data indicating the presence of a rhino given that a rhino is present is also very large P(*Rhino_track* = "True" | *Rhino_present* = "True") = 0.8000. Rhinos tend to hide, thus there could be no tracking data while a rhino is indeed present (P(*Rhino_track* = "False" | *Rhino_present* = "True") = 0.2000).

6.4.7 Rhino sighted

The *Rhino_sighted* node communicates whether a rhino was sighted by tourists, rangers, or other individuals in a specific cell [25]. This node has two states corresponding to whether or not a rhino was sighted ("True" and "False").

Despite their size, rhinos are very adept at hiding in tall grass or under shrubs or bushes. Since a rhino cannot always be observed, we have to rely on other elements to inform us of the presence of rhinos. If a rhino is present at a certain location, it is not to say that the rhino will be observed. The same goes for a rhino that is not present at a certain location: tourists frequently mistake animals for tree trunks or piles of rocks.

External reports assist in locating rhinos. This is hard evidence, a yes or no answer obtained by asking rangers in a specific cell whether rhinos are present. The case where a rhino is present will increase the likelihood of a rhino being sighted ($P(Rhino\_sighted$ = "True" | $Rhino\_present$ = "True") = 0.2000). The likelihood of observing a rhino when there is a rhino present is quite small as rhinos like to hide in the tall grass. This likelihood, however, is still higher than observing a rhino when a rhino is not present ($P(Rhino\_sighted$ = "True" | $Rhino\_present$ = "False") = 0.1000). The likelihood of not observing a rhino when a rhino is not there is very high ($P(Rhino\_sighted$ = "False" | $Rhino\_present$ = "False") = 0.9000), as is the likelihood of a rhino being present, undetected ($P(Rhino\_sighted$ = "False" | $Rhino\_present$ = "True") = 0.8000). Table 6.24 illustrates these probabilities.

Table 6.24. *Rhino_sighted* CPT

States		
Rhino_present	*Rhino_sighted*	Probabilities
False	False	0.9000
True	False	0.8000
False	True	0.1000
True	True	0.2000

6.5 RANGER SUBGROUP

The ranger subgroup contains the available knowledge pertaining to the effectiveness of the rangers and their safeguarding procedures. The subgroup is presented in Figure 6.10.

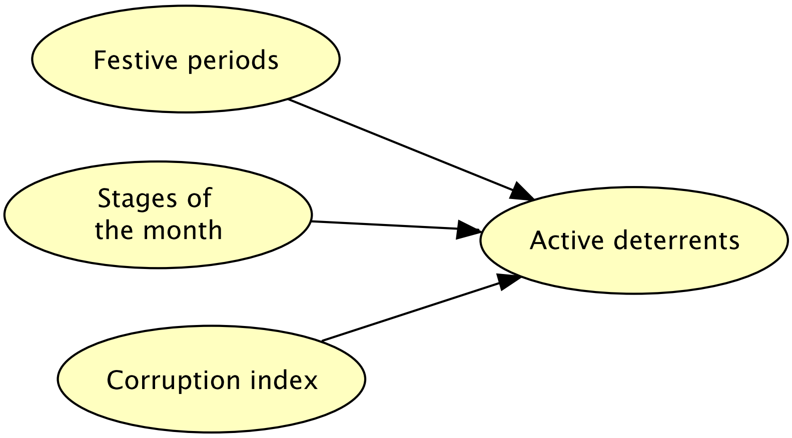

Figure 6.10. Ranger subgroup

6.5.1 Corruption index

The absence of a ranger plus the presence of both a rhino and a poacher increases the probability of a poaching event occurring, as per RAT. The *Corruption_index* node is a variable that conveys how corrupt a ranger or guardian is. If a ranger or other KNP employee is corrupt, it will be easy for poachers to bribe them to either help with the poaching attack, or to turn a blind eye. The corrupt rangers are either threatened or offered some kind of reward (usually monetary) for their cooperation.

During later discussions it became clear that "Corruption index" is not the best title for this variable, as corruption is not the only reason a ranger might be incapable. Other factors such as experience level, leave taken, exhaustion, or ill-discipline might be the cause for a ranger's incapability. The reader is referred to Section 8.4 for a discussion on this.

The states for this variable are "Corrupt" and "Not corrupt". There are no data available for this variable, yet it is known that there is a measure of corruption present. It was decided to assign a small weight to the "Corrupt" state. Table 6.30 shows the prior probabilities for *Corruption_index*.

Table 6.25. *Corruption_index* CPT

States	Probability
Not corrupt	0.9000
Corrupt	0.1000

6.5.2 Stages of the month

The timing of poaching events appears to be linked to the stages of the month, especially to payday. During times close to payday people are more regularly recruited for poaching, and as a consequence the effectiveness of deterrents is decreased. Rangers might also have to go to Skukuza to sign documents and therefore leave their posts. For a few days after payday, rangers may possibly visit shebeens and drink, divulging sensitive information regarding the rhinos or the safeguarding of the rhinos. The *Stages_of_the_month* node influences whether or not the rangers (as well as other deterrents) are effective in hindering poachers.

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. The states for the *Stages_of_the_month* node are "Payday - 7 days", "Payday + 7 days", "Payday + 14 days", and "Rest of the month". Assume that pd = Payday. The four state intervals can thus be written as $[pd-7, pd), [pd, pd+7), [pd+7, pd+14)$, and $[pd+14, pd-7)$. Table 6.26 illustrates the prior probabilities. The first three states have the same weight because each of these states contain seven days out of a possible 30 (the average number of days per month were taken to be 30). The fourth state has a slightly higher probability as it has to cover a nine day interval.

Table 6.26. *Stages_of_the_month* CPT

States	Probability
$[pd-7, pd)$	0.2300
$[pd, pd+7)$	0.2300
$[pd+7, pd+14)$	0.2300
$[pd+14, pd-7)$	0.3100

6.5.3 Festive periods

After consulting with experts, it became clear that there are certain spikes of rhino poaching events during the year. These spikes seem to correspond to certain holidays and festive periods, not only in South Africa, but also Mozambique and the East. A graph would have been shown if it were not for the sensitivity of the matter.

The $Festive\_periods$ node influences the effectiveness of the deterrents. During festive periods, rangers might be on holiday, thereby decreasing the number of resources in the KNP. The states of $Festive\_periods$ are "False" and "True" depending on whether the specific date falls in a festive period or not.

A list was made of all the main holidays in South Africa, Mozambique, and China (such as Easter, Christmas, New Year, and Chinese New Year). The reason for including South African holidays on the list is self-explanatory. Mozambican public holidays were included as many poachers are from Mozambique, as are rangers. Chinese New Year was included as the demand during this period is one of the main driving forces behind the poaching epidemic. Rich and powerful Chinese individuals want to showcase their new rhino horns during festive periods, or give them as gifts. The increased tourist numbers during holidays also provide cover for poachers who want to enter the park, since the gates are very busy.

An arbitrary calendar year was used to count how many days in a year fall in these festive periods. The link between poaching events and festive periods is still largely unknown, thus it was decided to not only include the main festival periods, but also the week before and the week after each festive period. The ratio of festive periods to number of days in the year was then calculated. The resulting prior probabilities are illustrated in Table 6.27.

Table 6.27. $Festive\_periods$ CPT

States	Probability
False	0.6767
True	0.3233

6.5.4 Active deterrents

This variable used to be called "Ranger present" and would indicate the presence of a ranger in that cell. The *Active_deterrents* node indicates whether or not deterrent efforts are effective in deterring poachers. Active deterrents are specific to anti-poaching actions and are the elements put in place to stop rhino poaching such as ranger (patrols), visible sensors, alarms, and lights. The states for *Active_deterrents* are "Not effective" and "Effective".

Active_deterrents is also the combination of abstract factors relating to the effectiveness of deterrents, ranger corruption (*Corruption_index*), proximity to payday (*Stages_of_the_month*), and how close it is to festive periods (*Festive_periods*). Table 6.28 shows the probabilities for *Active_deterrents* given its parent nodes. This CPT has a total of $2 \times 4 \times 2 \times 2 = 32$ entries corresponding to the number of states of each node, thus only a few entries are shown in Table 6.28. The representation of the probability of *Active_deterrents* conditioned on *Corruption_index*, *Stages_of_the_month*, and *Festive_periods* can be written as:

$$P(Active\_deterrents|Corruption\_index, Stages\_of\_the\_month, Festive\_periods).$$

Table 6.28. *Active_deterrents* CPT

Corruption_index	*Stages_of_ the_month*	*Festive_periods*	*Active_deterrents*	Probabilities
Not corrupt	Payday - 7	False	Not effective	0.5200
Corrupt	Payday - 7	False	Not effective	0.8080
Not corrupt	Payday + 7	False	Not effective	0.6000
Corrupt	Payday + 7	False	Not effective	0.8400
⋮	⋮	⋮	⋮	⋮
Corrupt	Rest of month	True	Effective	0.2240

The ideal situation for rangers to be effective would be when rangers are capable, the specific date is far from payday, and there are no immediate festive periods. Although this CPT is smaller than others, it is still large. The CPT was broken up into smaller tables to make elicitation easier as was done with *Accessibility*.

Table 6.29 shows the simplified CPT for the case where *Active_deterrents* = "Effective" and *Corruption_index* = "Not corrupt". Table 6.30 illustrates the simplified CPT for the case where *Active_deterrents* = "Effective" and *Corruption_index* = "Corrupt". *Active_deterrents* has two states, one of which is contained in each of these tables. The probabilities for the same configurations but for *Active_deterrents* = "Not effective" is the complement of the entries in Table 6.29 and Table 6.30.

Table 6.29. Simplified CPT for *Active_deterrents* = "Effective", *Corruption_index* = "Not corrupt"

Festive_periods		*Stages_of_the_month*			
		Payday - 7	Payday + 7	Payday + 14	Rest of month
		0.6000	*0.5000*	*0.8000*	*1.0000*
False	*1.0000*	0.4800	0.4000	0.6400	**0.8000**
True	*0.7000*	0.3360	0.2800	0.4480	0.5600
Corruption_index.Not corrupt					

Table 6.30. Simplified CPT for *Active_deterrents* = "Effective", *Corruption_index* = "Corrupt"

Festive_periods	*Stages_of_the_month*			
	Payday - 7	Payday + 7	Payday + 14	Rest of month
False	0.1920	0.1600	0.2560	0.3200
True	0.1344	0.1120	0.1792	0.2240
Corruption_index.Corrupt				

6.6 HISTORICAL VULNERABILITY

The *Historical_vulnerability* node is excluded from the latest model, but it is important to discuss the work that was performed and the reasons for excluding it from the current model. This node fed directly into *Poaching_event* and summarised the possible effect that past poaching incidents has on future poaching incidents. The *Historical_vulnerability* node conveyed the measure of vulnerability assigned to a specific cell. The vulnerability is how vulnerable a cell is for poaching events given the history of poaching events in that cell. This is the only node that would have been populated with poaching data.

Assume that cell i is the cell under consideration. Studying the history of poaching events across the KNP for the past few years have shown that cell i has had five poaching events. The cell next to cell i has had zero poaching events during the same timeframe. It can thus be deduced that cell i is more likely to be the location for a future poaching event when compared to the cell next to it. All the poaching events per cell are counted and the ratio for each cell is calculated according to the total number of poaching events. If a cell's ratio is less than a certain threshold, the historical vulnerability for that cell is said to be low, conversely, if it is higher than the threshold it is said to be high. If the cell's ratio is high, it gives the cell a prior of being a possible hotspot for poaching attacks in future.

It was decided to exclude $Historical\_vulnerability$ from the model because we wanted to keep the model free from data in order to have a purely expert-driven model. This could later be compared to a purely data-driven model which is currently the focus of another study. The ideal would be to use both models and create a "best of both worlds" model where both expert knowledge and data is used. It was also thought that $Historical\_vulnerability$ contains information that is already incorporated elsewhere in the model.

An interesting feature of this node was the factors that contribute to a poaching event. Different variables and features were compared to detect patterns from poaching event data. A sensitivity analysis was performed to establish which features were the most important in informing us about the problem. It became evident that there were a handful of variables where one state seemed to dominate the rest of the states when analysing the data. An example of this is the moon illumination percentage where the full moon phase dominated the other phases.

The most obvious way of determining if a variable has any influence on poaching events is to count how many poaching events occurred in each of the states. The state with the most poaching events will then be the state with the most influence on poaching events. This, however, is not the case. Assume that 1000 poaching events occurred in the KNP and that 20 of them occurred within two kilometres of gates (close). The fraction of the poaching events can be calculated as 20/1000 = 2%. The 2% is the weighted variable and 20 is the "raw" variable. Assume now as a second part to the example that 2000 poaching events occurred in the KNP and that 30 of them occurred further than five kilometres from gates (far). The fraction of poaching events occurring far from gates is thus 30/2000 = 1.5%. More rhinos are poached far from gates when compared to those close to gates (30 versus 20). If only the unweighted “raw” data or variables are considered, it will be incorrectly concluded that rhinos are more at risk far from gates as opposed to close to gates. If the weighted variables are considered it can be concluded that rhinos are actually more at risk close to gates (20/1000 = 2% close to gates versus 30/2000 = 1.5% far from gates). This is the method that was followed to ascertain which states are more important for poaching attacks.

Figures 6.11-6.13 illustrate that $Distance\_to\_incursion\_points$ = “In neighbouring cell”, $Proximity\_to\_static\_deterrents$ = “Far from cell”, $Vegetation\_density$ = “Impenetrable vegetation”, $Landscape\_preference$ = “Neutral”, and $Moon\_illumination\_percentage$ = “75% - 100%” are the only variables to show a significant spike for a single state, thus it can be hypothesised that these variables have a significant influence on the likelihood of poaching events. The specific sequence of nodes in certain states leading to a poaching event can unfortunately not be deduced from the histograms. The other variables that do not contain a spike seem to have a weaker influence on poaching events. Further experiments would have to be performed to verify how much these variables contribute to poaching events.

Policy changes could also have an effect on the variables, but it is impossible to measure as it is not known when these policy changes occurred. For instance, the full moon was a very important indicator of poaching activity a year ago, but then the rangers were sent to certain hotspots every full moon, and so the poachers changed their operations. The data would thus reflect that the full moon was not a significant contributing factor. It is not known when these policy changes occurred thus it is impossible to say which variables were more important for a certain time period.

The mathematical calculation usually performed to compute error bars involve the mean, but since count data are used in this node, only the fraction of poaching events to cell counts are used. Table 6.31 illustrates a few important variables for computing the error bars.

It follows from the above that $\sum ccounts_i = N_c$ and $\sum pcounts_i = N_p$. The weights W_i are calculated as the fraction of cells falling in a state, and h_i is the fraction of poaching events to cells for a specific state i. The equation for calculating the variance of the sum of the weights S_k is given by Equation 6.4 which calculates how far each weight is from the mean:

Table 6.31. Variables for Calculating the Error Bars

N_p	Total number of poaching events
N_c	Total number of cells in the grid
i	Number of states in the particular node
$ccounts_i$	Number of cells falling in a specific state
$pcounts_i$	Number of poaching events occurring in a specific state
W_i	Weights (one for each state)
S_k	Sum of all weights of the events per bin k
h_i	Number of poaching events per state divided by the number of cells per state

$$\mathrm{var}[S_k] = \sum W_i^2 - \frac{1}{N} \sum (W_i)^2. \tag{6.4}$$

The standard deviation is used to construct the histograms, thus the square root of Equation 6.4 is used. In this study the weights are the fraction of cell counts, and summing the weights corresponds to multiplying the weights with the number of poaching counts for that state. All the histograms in this study are drawn with error bars of one standard deviation.

Figure 6.11 illustrates the marginal histograms for $Distance\_to\_incursion\_points$ and $Proximity\_to\_static\_deterrents$. The marginal histogram for $Distance\_to\_incursion\_points$ show that poachers prefer to poach where an incursion point is close, but more or less five kilometres away. They also prefer to poach far from static deterrents. In both of these variables the states are statistically significantly different, because their error bars do not overlap.

Figure 6.12 presents the marginal histograms for $Vegetation\_density$, $Landscape\_preference$, $Water$, and $Festive\_periods$. From these four histograms it appears that poachers prefer to poach in very dense vegetation, in landscape that is neutral to rhinos, and during festive periods. In all four these variables the states of the respective variables are statistically significantly different, because their error bars do not overlap.

Lastly, Figure 6.13 illustrates the marginal histograms for $Stages\_of\_the\_month$ and $Moon\_illumination\_percentage$. It appears that there are not any distinct patterns to observe

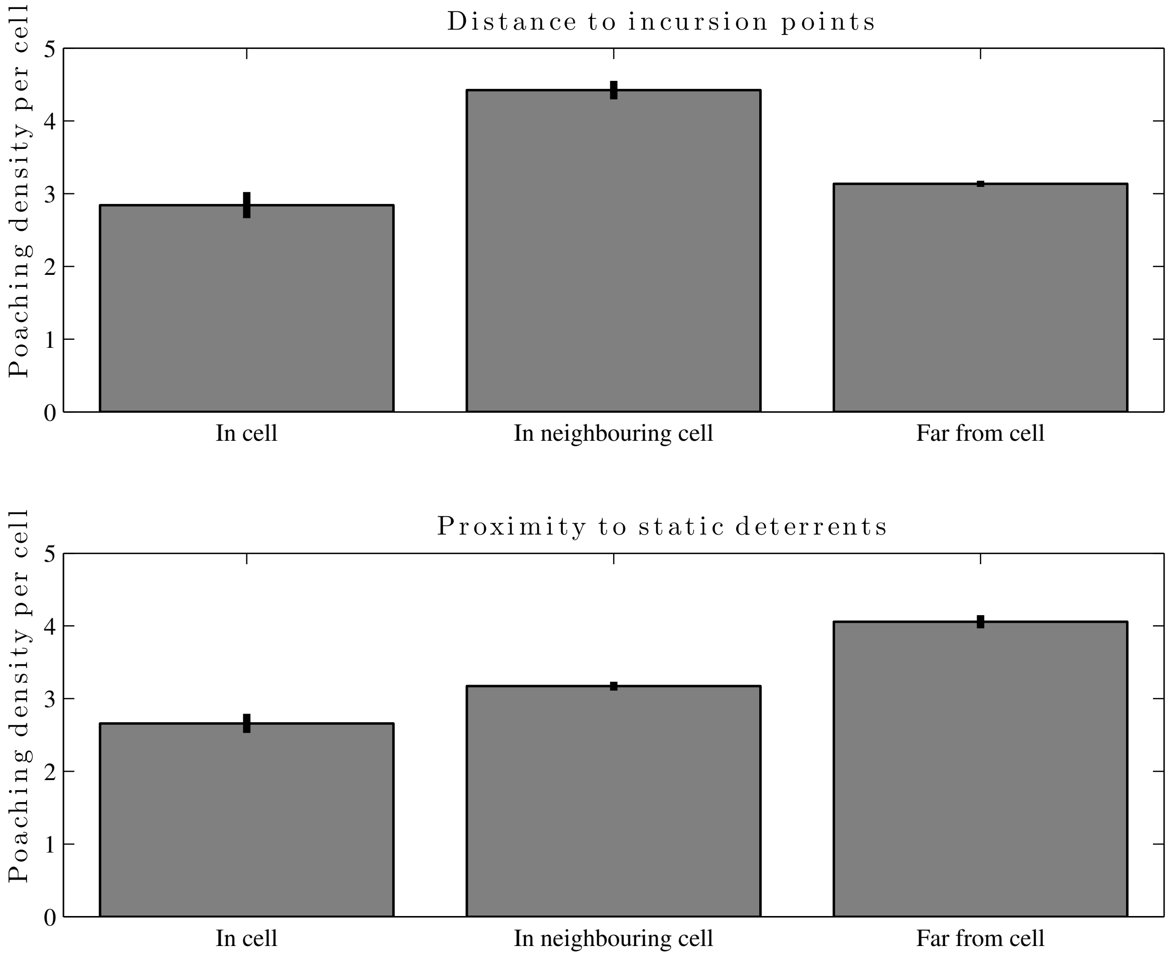

Figure 6.11. Marginal histograms

in these two variables. The error bars overlap for the middle two states of *Stages_of_the_month*, thus it can be said that these states are not statistically significantly different. The first and fourth states are, however, statistically significantly different. The $0-25\%$ and $25-50\%$ states of *Moon_illumination_percentage* are also not statistically significantly different, but the $50-75\%$ and $75-100\%$ states are statistically significantly different, because their error bars do not overlap.

6.7 POACHING EVENT SUBGROUP

The poaching event subgroup is the core of the model and contains all the available knowledge pertaining to whether or not a poaching event is going to occur. The subgroup is illustrated in Figure 6.14.

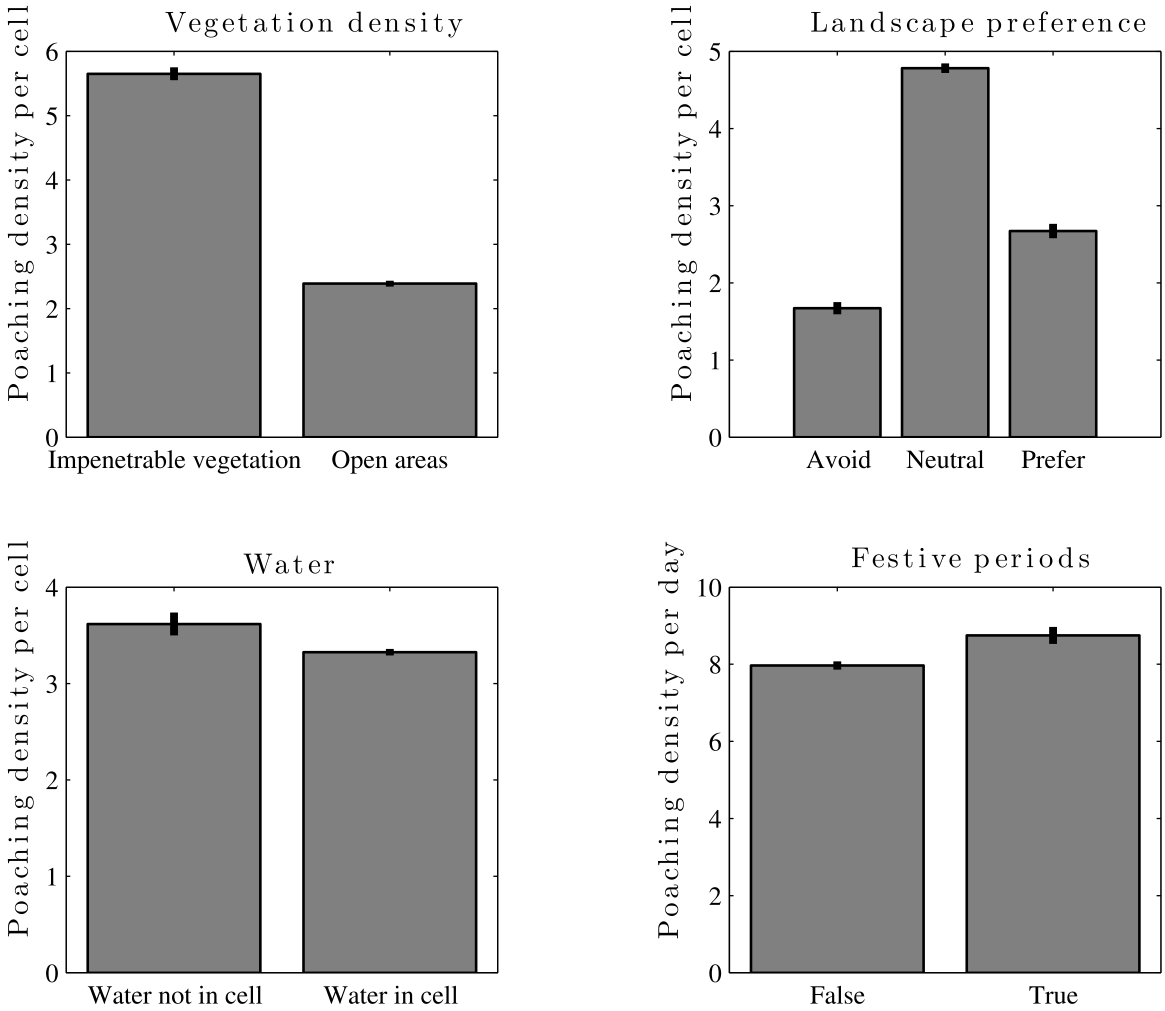

Figure 6.12. Marginal histograms

6.7.1 Poaching event

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

Poaching_event is the culmination of factors leading to a poaching event. These factors are concerning the effectiveness of rangers and other deterrents (*Active_deterrents*), if a poacher is present (*Poacher_present*), and if a rhino is present (*Rhino_present*). Table 6.32 shows the prior CPT for this node given its parent nodes. This CPT has a total of $2 \times 2 \times 2 \times 2 = 16$ entries corresponding to the number of states per node. The representation of the probability of *Poaching_event* conditioned on its

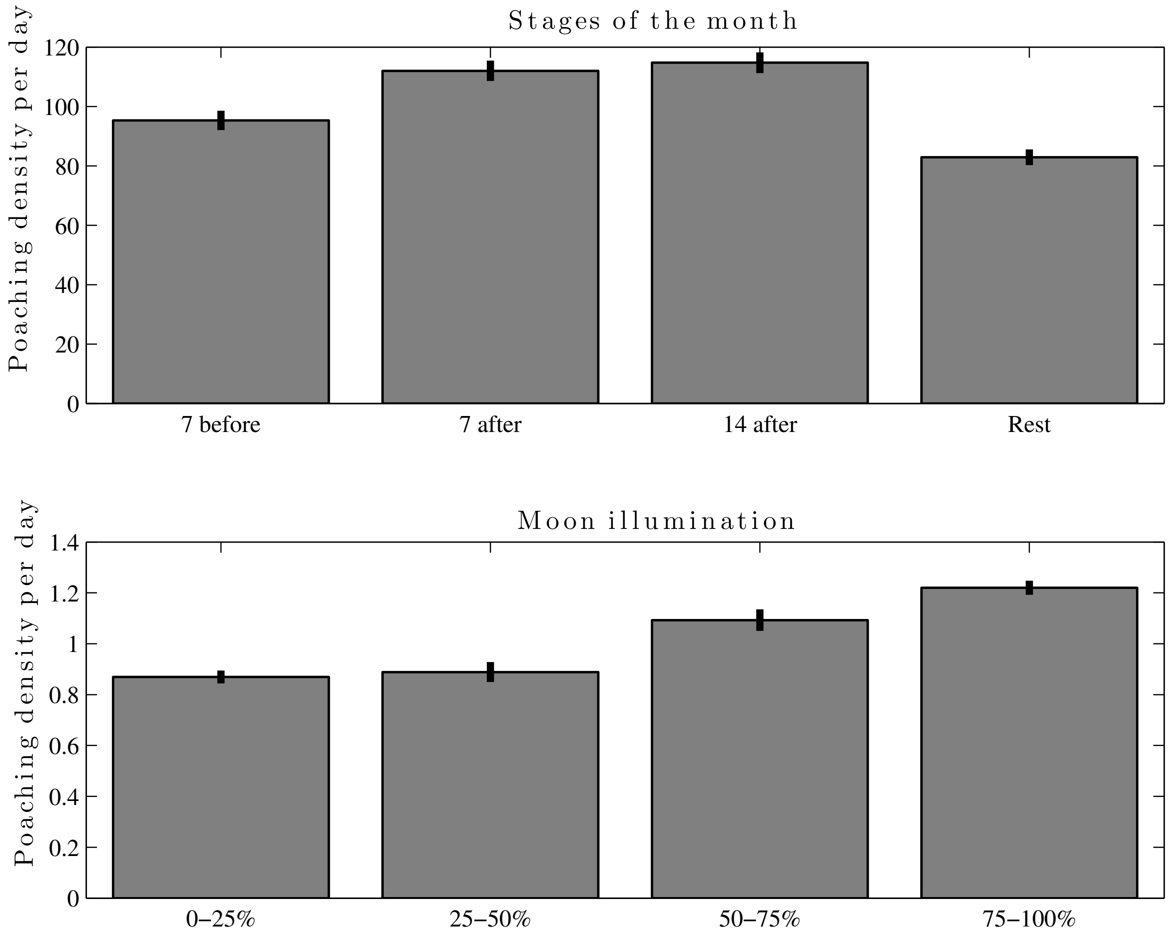

Figure 6.13. Marginal histograms

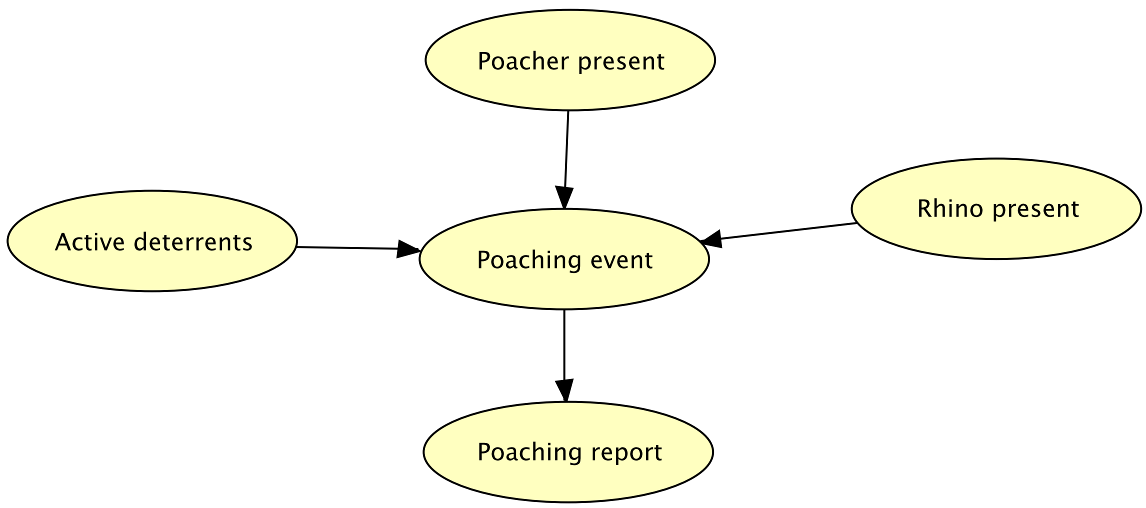

Figure 6.14. Poaching subgroup

parent nodes can be written as:

$$P(Poaching\_event|Poacher\_present, Rhino\_present, Active\_deterrents).$$

Table 6.32. *Poaching_event* CPT

States				
Poacher_ present	*Rhino_ present*	*Active_ deterrents*	*Poaching_ event*	Probabilities
False	False	Not effective	False	1.0000
True	False	Not effective	False	1.0000
⋮	⋮	⋮	⋮	⋮
False	False	Not effective	True	0.0000
True	False	Not effective	True	0.0000
False	True	Not effective	True	0.0000
True	True	Not effective	True	0.9800
False	False	Effective	True	0.0000
True	False	Effective	True	0.0000
False	True	Effective	True	0.0000
True	True	Effective	True	0.2940

Table 6.33 represents the simplified CPT for *Poacher_present* = "True" and *Active_deterrents* = "Not effective". A poaching event can only occur if there is both a rhino and a poacher present, thus only one entry in the table is non-zero. Table 6.34 illustrates the simplified CPT for *Poacher_present* = "True" and *Active_deterrents* = "Effective". The weighting between the tables is 0.3000.

Table 6.33. Simplified CPT for *Poacher_present* = "True", *Active_deterrents* = "Not effective"

		Poacher_present	
Rhino_present		False	True
		0.0000	*1.0000*
False	*0.0000*	0.0000	0.0000
True	*1.0000*	0.0000	**0.9800**
Active_deterrents.Not effective			

Table 6.34. Simplified CPT for *Poacher_present* = "True", *Active_deterrents* = "Effective"

	Poacher_present	
Rhino_present	False	True
False	0.0000	0.0000
True	0.0000	0.2940
Active_deterrents.Effective		

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 (here 0.9800). A poaching event cannot take place if a rhino or a poacher is not present, thus many of the entries in the table are zero.

6.7.2 Poaching report

Poaching_report is a binary variable that takes on the value of "True" if a poaching report is generated, and "False" otherwise. If a poaching event occurred, the probability of a poaching report would be high, unless there is corruption involved (in which case a ranger might want 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. Table 6.35 shows the prior probabilities for *Poaching_report*.

Table 6.35. *Poaching_report* CPT

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

6.8 GAINING A BETTER PERSPECTIVE

The previous sections dealt with describing and explaining the details of the model. The model consisted of subgroups, and each subgroup consisted of nodes. Subgroups simplify the process of obtaining probabilities and also of understanding the model.

The model in Figure 6.15 is obtained by putting all the puzzle pieces (the subgroups and their nodes and edges) together. This is the same model as was presented in Figure 6.2, but now without an explicit partitioning of subgroups. This model represents all that is known about the rhino poaching problem in the KNP.

The relationships between the nodes in a BN can be written as probability distributions. The joint probability distribution for the model in Figure 6.15 is shown by Equation 6.5. The probability distributions are used to perform inferences on the model and to answer questions about the model. This will be described in detail in the following chapters. Table 6.36 provides the reader with a list of variable names for Equation 6.5.

$$\begin{aligned} P(\Psi) &= P(\text{PR}|\ \text{PE})P(\text{PE}|\ \text{AD, PP, RP})P(\text{AD}|\ \text{CI, SOTM, FP})P(\text{CI})P(\text{SOTM}) \\ &\times P(\text{FP})P(\text{PPO}|\text{PP})P(\text{PP}|\ \text{A, TOD, MIP, We})P(\text{A}|\ \text{VD, DTIP, PTSD})P(\text{VD}) \\ &\times P(\text{DTIP})P(\text{PTSD})P(\text{TOD})P(\text{MIP})P(\text{We})P(\text{RP}|\text{LP, Wa, HRP})P(\text{LP}) \\ &\times P(\text{Wa})P(\text{HRP})P(\text{RT}|\text{RP})P(\text{RS}|\text{RP}), \end{aligned} \tag{6.5}$$

where Ψ denotes all the variables contained in the model.

Table 6.36. Variable names

Variable	Variable name
Vegetation_density	VD
Distance_to_incursion_points	DTIP
Proximity_to_static_deterrents	PTSD
Accessibility	A
Time_of_day	TOD
Moon_illumination_percentage	MIP
Weather	We
Poacher_present	PP
Positive_poacher_observation	PPO
Landscape_preference	LP
Water	Wa
Historical_rhino_presence	HRP
Rhino_present	RP
Rhino_track	RT
Rhino_sighted	RS
Corruption_index	CI
Stages_of_the_month	SOTM
Festive_periods	FP
Active_deterrents	AD
Poaching_event	PE
Poaching_report	PR

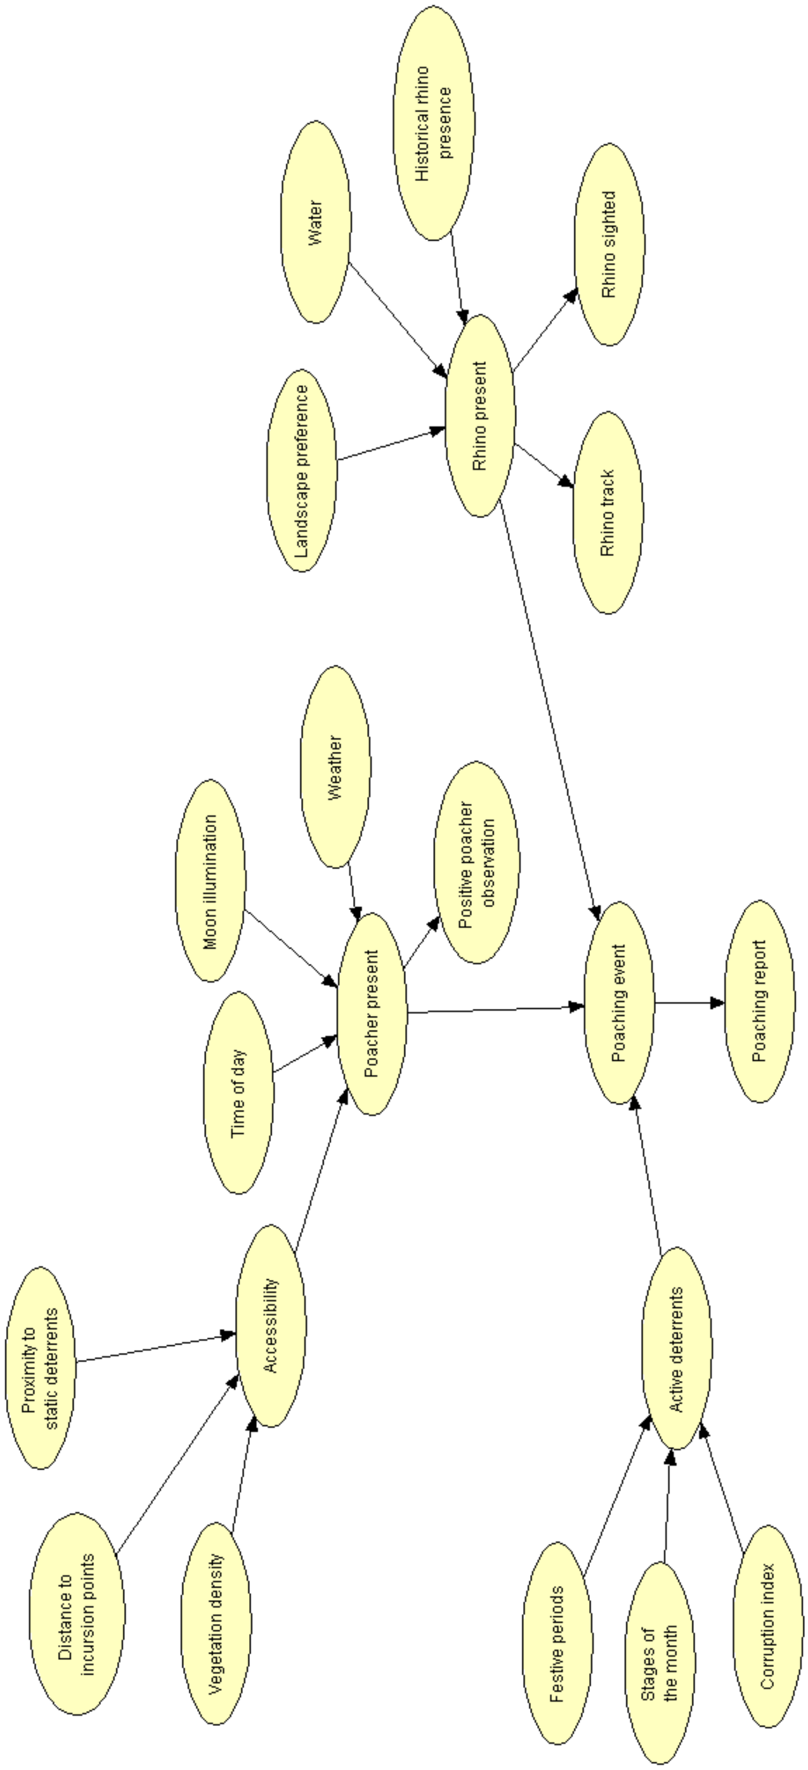

Figure 6.15. Rhino poaching model

6.9 EXPERIMENTING WITH DIFFERENT GRID SIZES

After the workshop certain experts asked what the result would be if the cell size of the grid was made finer. In order to inspect the difference it was decided to additionally investigate the effect of reducing the cell size to 1×1 kilometre.

Upon inspection it became clear that certain variables would have to be processed differently and the states of certain variables would have to change. For a cell size of 5×5 kilometres, the experts recommended states of "In cell", "In neighbouring cell", and "Far from cell" for both *Distance_to_incursion_points* and *Proximity_to_static_deterrents*. For a cell size of 1×1 kilometre, it was decided to change the states of *Distance_to_incursion_points* to "0 - 2.5 km", "2.5 - 5 km", and "> 5 km", where the distance was now calculated from the midpoint of the cell to the closest incursion point. Since the incursion points occur mostly on the border of the KNP, it was decided to not use one-kilometre intervals for the states, but rather two and a half kilometres. The states for *Proximity_to_static_deterrents* was changed to "0 - 5 km", "5 - 10 km", and "> 10 km" as the experts were convinced that poachers do not wish to be close to static deterrents and that this distance is not conditioned on the cell size. The *Water* variable was also altered and its states changed to "0 - 5 km" and "> 5 km" as that was the states suggested by the experts, even though the grid size is now 1×1 kilometre. This variable was also processed by computing the distance from a cell's midpoint to the closest water point.

6.10 CONCLUSION

In this chapter the architecture of the BN model is presented after it was altered and validated by experts. The model was divided into four logical subgroups for ease of viewing and elicitation, coinciding with RAT stating the elements that need to be in place for a crime to occur. All the explanations and motivations were at the hand of the subgroups: poacher subgroup, rhino subgroup, ranger subgroup, and poaching subgroup. The choice of nodes and states were discussed and motivated, as well as how the probabilities were derived and computed.

Partitioning the model into parts added a certain richness as the interaction between the different parts could be inspected and studied. The subgroups make it possible to present different aspects of the model to relevant experts. This increases the confidence in each of the subgroups, as well as in the totality of the model. The causal nature of the model can be viewed as the pieces of a puzzle - each an important element but fitting together to form a complete picture. The experts concur that the main building blocks of the model correspond to their own beliefs and experience, and that the synthesis of the building blocks also corresponds to their knowledge.

The effect of reducing the cell size to 1×1 kilometre was investigated and a number of variables' states had to be re-processed to better correspond with the finer grid. The outcome of these changes are discussed in Chapter 8.

The model presented in this chapter serves as a more realistic view on the current poaching problem and one that can easily be updated and used in a anti-poaching strategy. It is important to note that no other causal model currently exists for predicting the rhino poaching problem. The following chapter describes a software application to the rhino poaching problem and showcases templates that were built specifically for expert elicitation.

CHAPTER 7 RHINO POACHING APPLICATION

7.1 INTRODUCTION

The previous chapter described the latest model, from the subgroups to the individual nodes. The Conditional Probability Tables (CPTs) were populated with expert knowledge due to a lack of complete rhino poaching data. This chapter describes a template that was used to simplify expert elicitation as well as easily create a Bayesian Network (BN) structure.

The work for this thesis was performed mainly in Matlab® (from here on referred to as "Matlab") [98] as Matlab is an environment ideally suited to matrix computing. Hugin® (from here on referred to as "Hugin") [99] is a popular tool for BNs, but it proved a very difficult task to integrate it with Matlab. Kevin Murphy's BN toolbox (BNT) [91] was used which is a Matlab-based toolbox for BNs. The downside is that a ranger will most likely not have access to Matlab. An alternative to Matlab® and Hugin® is to use open-source software such as R.

Visualising a BN in the BNT is challenging and not visually pleasing. Every time a node or an arc is changed in the BNT (moved, added, or removed), the CPTs contained by that node also change. In the BNT that often signifies that the entries of a very large matrix has to be retyped. It is frustrating not being able to see if you are inputting the probabilities at the correct position in the matrix, consequently there is a chance for errors.

These problems were addressed by developing and maintaining the structure of the BN in Microsoft Excel® (from here on referred to as just "Excel"). The author created a template where the user has to enter details about the BN after which it is imported into Matlab. Many people, especially companies, have access to Excel and this was one of the main reasons behind the decision to use Excel. Another reason for using Excel is that tables and matrices can be visualised effectively. The BN structure can thus be created in Excel but imported into Matlab to create the BN model.

7.2 BUILDING A BAYESIAN NETWORK IN MATLAB®

This section describes the BNT and how to use it, after which the Excel template is illustrated for the development and maintenance of the BN. The example used is a simple one from the BNT manual that was adapted from Russell and Norvig [100]. The reason for using this non-rhino example is for ease of reference if the reader would like to go back to the example in Murphy [91] and replicate it.

7.2.1 Development

A man walks outside one morning and sees that the grass is wet (*Grass_wet*). It could have been that it had been raining (*Rain*, *Cloudy*), or it could have been that the sprinkler was on during the early hours of the morning (*Sprinkler*). Certain assumptions can be made about this situation. If it was cloudy, then there is a good chance for rain. If it had rained, it definitely had to have been cloudy at some point. Also, if it had rained there would have been no need for the sprinkler to go on. These assumptions are the causal links between the nodes.

The sprinkler network consists of four nodes (*Cloudy*, *Sprinkler*, *Rain*, *Grass_wet*) and four edges. *Cloudy* has two edges, one pointing to *Sprinkler* and one pointing to *Rain*. *Sprinkler* has an edge pointing to *Grass_wet* and *Rain* also has an edge pointing to *Grass_wet*. Tables 7.1-7.4 illustrate the respective probabilities.

Table 7.1. *Cloudy* (1) probabilities

$P(Cloudy = False)$	$P(Cloudy = True)$
0.50	0.50

Specific probabilities can be assigned to different scenarios. The probability of the grass being wet given that the sprinkler was on, as well as it having rained, is $P(Grass\_wet = True|Sprinkler = True, Rain = True) = 0.99$. The probability of the grass being dry given that the sprinkler was on, as well as it having

Table 7.2. *Sprinkler* (2) probabilities

	$P(Sprinkler = False)$	$P(Sprinkler = True)$
$Cloudy = False$	0.50	0.50
$Cloudy = True$	0.90	0.10

Table 7.3. *Rain* (3) probabilities

	$P(Rain = False)$	$P(Rain = True)$
$Cloudy = False$	0.80	0.20
$Cloudy = True$	0.20	0.80

Table 7.4. *Grass_wet* (4) probabilities

		$P(Grass\_wet = False)$	$P(Grass\_wet = True)$
$Sprinkler = False$	$Rain = False$	1.00	0.00
$Sprinkler = True$	$Rain = False$	0.10	0.90
$Sprinkler = False$	$Rain = True$	0.10	0.90
$Sprinkler = True$	$Rain = True$	0.01	0.99

rained, is $P(Grass\_wet = False|Sprinkler = True, Rain = True) = 0.01$. These probability tables can be populated by expert knowledge, empirical data, or literature.

Now that the structure and the parameters are defined, the Directed Acyclic Graph (DAG) can be specified. The way to do this in Matlab is to create an adjacency matrix. An adjacency matrix is an $N \times N$ matrix (where N is the number of nodes in the network) that gives a mapping according to which nodes are connected to which other nodes. The diagonal of this adjacency matrix will always be zeros, as a node cannot connect to itself. An adjacency matrix is equivalent to an edge list and is a "machine-readable description of a graph" [42].

The nodes have to be numbered in topological order: ancestors before descendants. If there exists an edge (i, j) from node i to j, then $i < j$ must hold [91]. This means that node 5 → node 8, but not node 8 → node 5. There is no unique way of numbering the nodes, but several algorithms exist to compute a correct numbering scheme [101–105]. The four nodes in the sprinkler network was numbered as follows: *Cloudy* = 1, *Sprinkler* = 2, *Rain* = 3, *Grass_wet* = 4. In the rhino poaching BN the nodes were numbered from left to right, top to bottom, and clockwise (where possible).

The size and the type of each node also have to be defined. Nodes are either discrete or continuous. If a node is discrete, then the number of states will be the number of different values the node can take on. A continuous node can be defined as a vector where the length of the vector corresponds to its size. In the sprinkler example, all the nodes are discrete and binary (each node has two states). Next the observed nodes need to be defined, but if they are unknown the empty list is used.

The next step is to define the parameters. The parameters of a BN are represented by a Conditional Probability Distribution (CPD) that defines the relationship between a node and its parents. Only the nodes with an arc leading to the child node are called parent nodes; thus a child node can have multiple parent nodes, but they are all on the same tier in the network. The node in the network is the same as the variable in the problem, so the two names will be used interchangeably for the rest of this thesis. The simplest, and most-used, type of CPD is the table, called the CPT. The CPT is used when all the nodes involved are discrete. The important thing to remember when working with CPTs in Matlab is that the nodes should be used in their numbered order, with the child node last. Where software programs like Hugin lets you choose your own order (and presents the CPTs as a long horizontal matrix), the Matlab BNT forces you to use their convention (and presents the CPT as a long column).

The CPT for *Grass_wet* (4) will have *Sprinkler* (2), then *Rain* (3), and then *Grass_wet* (4) as columns, as the child node is always last. Matlab is indexed at one (unlike C that is indexed at zero), and its arrays are stored in memory such that the first element always switches fastest. Table 7.5 illustrates this concept for the CPT for *Grass_wet* (4). Note that a value of one is assigned to "False" and a value of two is assigned to "True". If a node does not have parents, it is called a root node, and its CPT is a column vector denoting its prior.

Depending on whether the node is discrete or continuous, and if it is a root node, there are different CPD constructors contained in the BNT toolbox for each type of node. The main constructors are `tabular_CPD`, `root_CPD`, `softmax_CPD`, and `gaussian_CPD` [91].

Table 7.5. *Grass_wet* (4) CPT

Sprinkler	*Rain*	*Grass_wet*	Probability
False (1)	*False* (1)	*False* (1)	1.0000
True (2)	*False* (1)	*False* (1)	0.1000
False (1)	*True* (2)	*False* (1)	0.1000
True (2)	*True* (2)	*False* (1)	0.0100
False (1)	*False* (1)	*True* (2)	0.0000
True (2)	*False* (1)	*True* (2)	0.9000
False (1)	*True* (2)	*True* (2)	0.9000
True (2)	*True* (2)	*True* (2)	0.9900

Tabular nodes

As already explained, if a node and all its parents are discrete, the CPD can be denoted by a table (CPT) or matrix [91].

Softmax nodes

If a node is discrete, but has one or more parents that are continuous, the CPD is specified by a softmax function. A softmax (or normalised exponential) function can be used to create a weighed average that represents a "smoothed version" of the max function [73] and acts like a soft thresholding operator [91]. The softmax function is a generalisation of the logistic function when there are only two categories.

Gaussian nodes

The Gaussian distribution is used for continuous-valued nodes. There are four different distributions that can be defined on a child node: no parents (root node), continuous parents, discrete parents, and continuous and discrete parents. The Gaussian node can be constructed with the `gaussian_CPD` constructor, either with random parameters or with specified values.

Root nodes

A root node can be discrete or continuous, but because the node does not have any parents, its CPD is represented by a column vector depicting its prior.

7.2.2 Inference

The two objectives of a BN are 1) to structure the problem, and 2) to answer questions about the problem. The first objective was met in the previous section, and in this section the objective is addressed of answering questions about the problem. In the sprinkler example many different questions could have been asked, such as what the probability is that the sprinkler had been on, or what the probability is that it had been raining. Reasoning about these types of questions is called *inference*. The BNT has many different algorithms (or engines) for inference that differ according to speed, complexity, and accuracy. The most popular and best-for-all-situations engine is the junction tree engine which, according to Murphy [91] is the "mother of all exact inference algorithms". The other engines have equivalent constructors, but differ in the parameters needed to execute the algorithms.

7.2.3 Marginal and joint distributions

In order to answer questions about the model, a marginal distribution has to be computed. Suppose the question being asked is, what is the probability of the sprinkler being on, given that the grass is wet? The grass is thus definitely wet (it is the evidence that is added: $Grass\_wet = 2$), and we would like to know what the probability is that the sprinkler was on. All the other nodes in the network are said to be unobserved or hidden. This evidence is added to the inference engine. The `marginal_nodes` function is used to compute the marginal distribution, which returns a vector containing the probabilities for each state. In this case, $P(Sprinkler = 2|Grass\_wet = 2) = 0.4298$. These probabilities were tested in Hugin.

Further evidence can be added, such as that it rained ($Rain = 2$). If the above experiment is executed again with the added evidence of rain, the computed probability is $P(Sprinkler = 2|Grass\_wet = 2, Rain = 2) = 0.1945$. This probability is lower than the previous probability because the fact that it rained decreases the chance of the sprinkler having been on.

The joint distribution of the nodes can be computed in roughly the same way as marginal distributions. No evidence is supplied, but the `marginal_nodes` function is supplied with an extra parameter containing the nodes we wish to compute the joint distribution on. When a node is sampled, it implies that all its parents are observed. When "Bayesian learning" is stated, it means that the posterior probability is calculated over the parameters given fully observed data [91].

7.3 APPLICATION TO THE RHINO POACHING PROBLEM

The same method is followed to develop the rhino poaching BN. An Excel workbook was created for the rhino poaching BN with sheets named *Nodes, States,* and *Relationships*. The sheet named *Nodes* contains a list of the topologically ordered nodes together with their numbers, the number of states per node, and the node's type (discrete or continuous). Figure 7.1 illustrates the *Nodes* sheet. This sheet provides the structure and the outline of the BN to be computed.

The *States* sheet contains once again the node numbers, node names, and each node's respective states. In the sprinkler example every node had two states namely "False" and "True", but in the rhino poaching BN the states are often not binary, and it is important to know how many states there are, as well as the names of the various states. The *Relationships* sheet contains the parent node number, parent node name, child node number, and child node name. Each row represents a single relationship, so if a child node C_i has two parent nodes, then there will be two rows where the node C_i is in the child node column. In Excel it is easier to work with long columns than long rows. The number of nodes and the number of arcs specify how long the table is, and the width of the table is fixed to parent node number, parent node name, child node number, and child node name.

In future applications these three sheets (*Nodes, States, Relationships*) will be completed by the user, and any changes made to the BN is made in this Excel workbook. The way the Excel workbook and the Matlab code was written was to ensure generality. This template can be used to create any BN as well as maintain its structure.

A Matlab script evaluates this workbook's first three sheets (*Nodes, States, Relationships*) and creates CPT outlines. Based on the first sheet (*Nodes*), the code also calculates which nodes are root nodes,

	A	B	C	D
1	Node Number	Node Name	Number of states	Type
2	1	Vegetation density	2	discrete
3	2	Distance to incursion points	3	discrete
4	3	Proximity to static deterrents	3	discrete
5	4	Accessibility	2	discrete
6	5	Time of day	3	discrete
7	6	Moon illumination	4	discrete
8	7	Weather	4	discrete
9	8	Poacher present	2	discrete
10	9	Positive poacher observation	2	discrete
11	10	Landscape preference	3	discrete
12	11	Water	2	discrete
13	12	Historical rhino presence	4	discrete
14	13	Rhino present	2	discrete
15	14	Rhino track	2	discrete
16	15	Rhino sighted	2	discrete
17	16	Corruption index	2	discrete
18	17	Stages of the month	4	discrete
19	18	Festive periods	2	discrete
20	19	Active deterrents	2	discrete
21	20	Poaching event	2	discrete
22	21	Poaching report	2	discrete

Figure 7.1. The *Nodes* sheet

which are continuous nodes, and which are softmax or Gaussian nodes. A fourth sheet, *CPT*, contains the newly created CPT outline (see Figure 7.2). The user then completes the CPT probabilities.

In the case of very large CPTs (for instance, if a node has three or more parents), smaller, more manageable CPTs are created in another workbook. These CPTs are easier for the expert to complete, and certain "shortcuts" can also be taken to complete CPTs. This is described in detail in Chapter 5 and Chapter 6.

The Matlab code can extract and process the newly entered probabilities from CPTs and assigns the probabilities to their correct node variables. These probabilities and nodes are then assigned to the correct constructors as mentioned in Subsection 7.2. Inference can then be performed on the network by entering evidence and choosing a node of interest. The model can currently be used to answer one

A	B	C	D	E	F
1	VegetationDensity				
	Impenetrable vegetation				
	Open areas				
2	DistanceToIncursionPoints				
	In cell				
	In neighbouring cell				
	Far				
3	ProximityToStaticDeterrents				
	In cell				
	In neighbouring cell				
	Far				
4	VegetationDensity	DistanceToIncursionPoints	ProximityToStaticDeterrents	Accessibility	
	Impenetrable vegetation	In cell	In cell	Low	
	Open areas	In cell	In cell	Low	
	Impenetrable vegetation	In neighbouring cell	In cell	Low	
	Open areas	In neighbouring cell	In cell	Low	
	Impenetrable vegetation	Far	In cell	Low	
	Open areas	Far	In cell	Low	
	Impenetrable vegetation	In cell	In neighbouring cell	Low	
	Open areas	In cell	In neighbouring cell	Low	
	Impenetrable vegetation	In neighbouring cell	In neighbouring cell	Low	
	Open areas	In neighbouring cell	In neighbouring cell	Low	
	Impenetrable vegetation	Far	In neighbouring cell	Low	
	Open areas	Far	In neighbouring cell	Low	
	Impenetrable vegetation	In cell	Far	Low	
	Open areas	In cell	Far	Low	
	Impenetrable vegetation	In neighbouring cell	Far	Low	
	Open areas	In neighbouring cell	Far	Low	
	Impenetrable vegetation	Far	Far	Low	
	Open areas	Far	Far	Low	
	Impenetrable vegetation	In cell	In cell	High	
	Open areas	In cell	In cell	High	
	Impenetrable vegetation	In neighbouring cell	In cell	High	
	Open areas	In neighbouring cell	In cell	High	
	Impenetrable vegetation	Far	In cell	High	
	Open areas	Far	In cell	High	

Figure 7.2. The *CPT* sheet containing the outline of the CPTs to be completed

of two questions: (1) Given a certain time frame, given a certain location, what is the probability of a poaching attack taking place? (2) What will the poaching situation look like over a certain time frame?

Populating the answer to the first question for multiple locations yields the second answer. Currently the size of a cell is 5×5 kilometres. To answer question (1), the user needs to specify a certain location in the form of GPS (Global Positioning System) coordinates, or by clicking on the map. A cell is then drawn around this location according to the cell size. Inferences can be drawn about what happens within this specific cell. Answering question (2) also uses a cell size of 5×5 kilometres, but this time the whole map is divided into cells of that cell size. All of this is done in Matlab with the help of the BNT and the Mapping Toolbox. The Mapping Toolbox is an additional toolbox that lets the user work with shapefiles and digital maps.

There are five steps in the process of obtaining a poaching probability if a query is lodged:

1. The user specifies a location (unless he/she wishes to populate the whole map).
2. The park is divided into cells of 5×5 kilometres and the location is locked.
3. The maps of the park are used as a lookup table to compute values of the specific cell for each of the nodes in the BN.
4. The BN is populated according to the values found in the previous step and the probability of a poaching event being true is calculated.
5. The probability of a poaching event is used to colour in the corresponding cell according to the probability.

7.4 USING THE BAYESIAN NETWORK

The process

There are three main components in the system, namely the user input module, the data pre-processing module, and the BN.

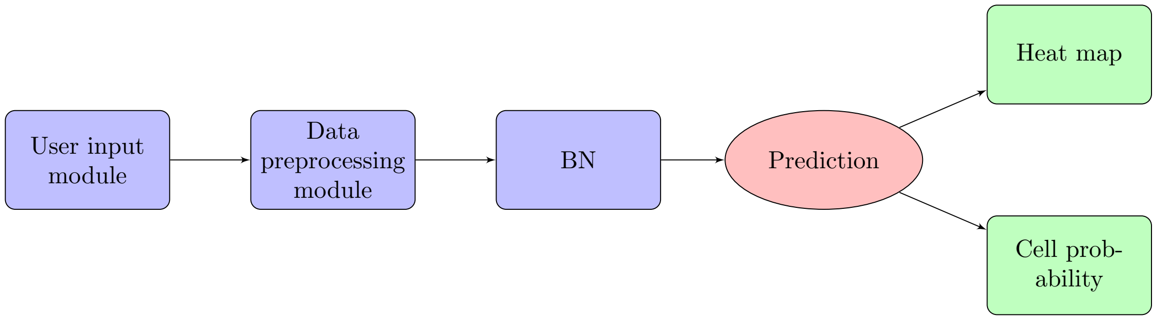

Figure 7.3. Main components of the system

The user input is an important part of the system, as this is the first part in calculating the necessary information for the BN. The user requests the probability of a poaching event, and the information he enters informs the result of the desired prediction. The pre-processing module extracts and processes the available information and populates the BN with the data.

The map of the Kruger National Park (KNP) is divided into cells. Predictions are made per cell, and each cell represents a summary of what happens in that cell, with regards to environment and terrain. Figures 7.4 and 7.5 illustrate the process in more detail. Figure 7.4 illustrates the process for a single cell, while Figure 7.5 illustrates the process for multiple cells, and in this case, all the cells in the grid. The map is divided into a grid of cells and then the user's input is requested. The map prior is calculated for a certain cell and the value for each node in the BN is populated. A probability for *Poaching_event* is computed, and is sent back to the user.

Each cell can be seen as a summary of all that is happening at that location in the KNP. Two cells will not have the same distance to water points, nor will they also have the same distance to variables such as incursion points. Each cell will thus have different states (evidence) for some of the nodes of the BN and that is why each cell is an instance of the BN. Populating (or providing evidence to) N cells in the KNP would mean that the BN would have to be populated and executed N times. Each time the BN is executed the states of the nodes for which evidence exist are determined. The configurations could be the same for different cells, but adjacent cells could also have entirely different configurations.

The user wants to know, given a certain time epoch, a certain date, and a certain geographical area, what the probability of a poaching event is. The user enters a date and chooses a time of day ("Midday", "Twilight", or "Night") from a list (corresponding to the three states of the $Time\_of\_day$ node), and then clicks on the map to select a location or specific area of interest. A grid is drawn over the map to guide the user in selecting an area of interest. The default size of the grid cells is 5×5 kilometres per cell. In an area the size of the KNP, this amounts to 954 possible 5×5 kilometres cells that are (at least partially) inside the KNP. The selected date, time of day, and coordinates of the area of interest are then sent to the pre-processing module.

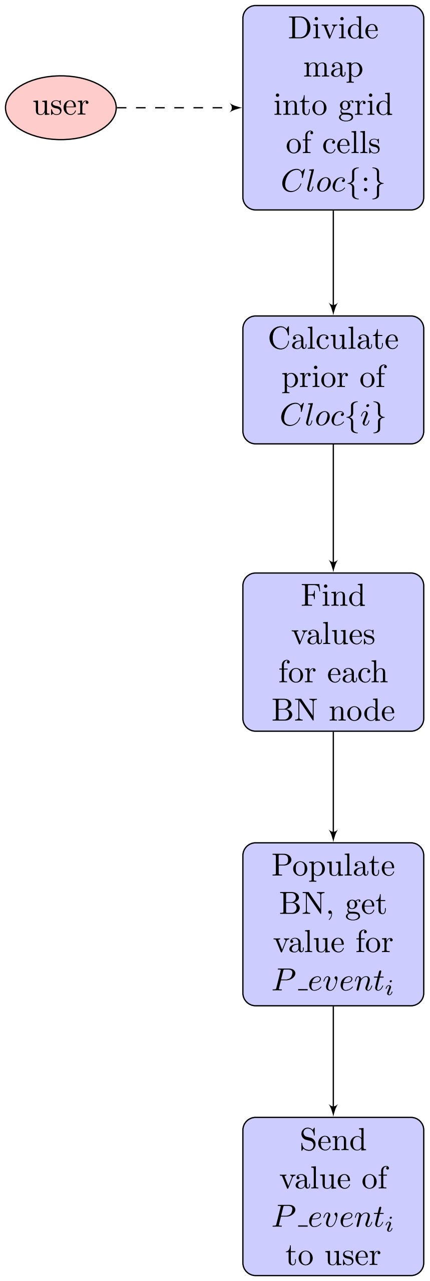

Figure 7.4. Process flow for a single cell

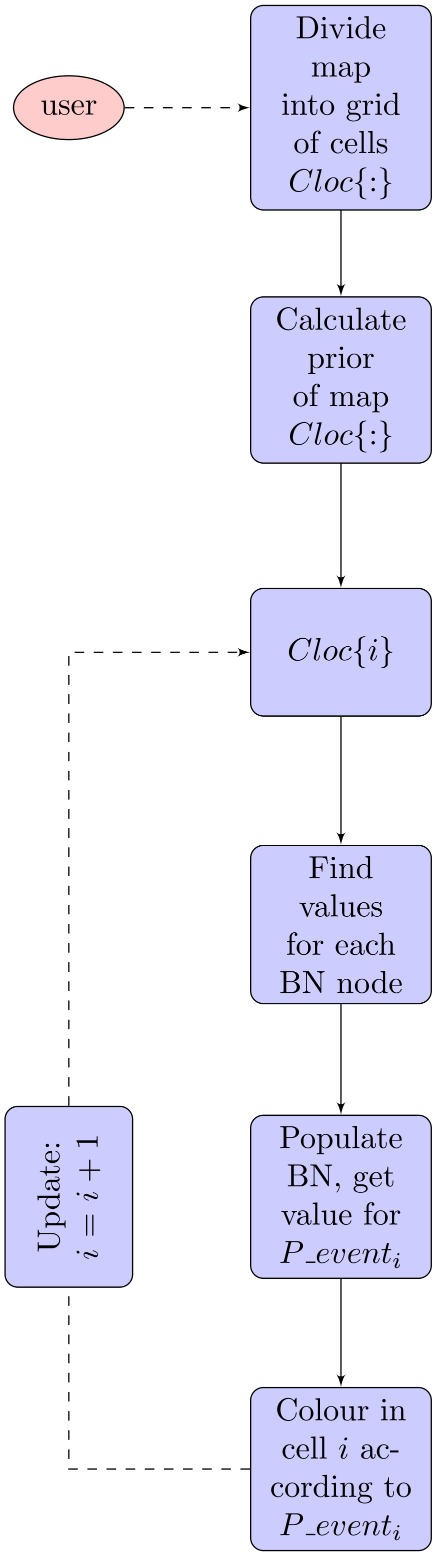

Figure 7.5. Process flow for all the cells

Pre-processing module

The pre-processing module uses the coordinates and determines the cell in which the chosen location falls. It does not matter where in the cell the user clicks (in the middle, on the edge, *etcetera*), the algorithm will determine that cell. The time of day entered is converted into a state. Each option presented to the user corresponds to a state of the $Time\_of\_day$ node. If the user selects "Midday", then the state of $Time\_of\_day$ will be $Time\_of\_day = 1$, because that is the value corresponding to "Midday".

The specific cell index is calculated from the coordinates where the user clicked. The cell index informs the pre-processing module where to find the evidence it needs. Knowing the cell index, the processing module can look up the states of all the static nodes such as $Water$ and $Landscape\_preference$. Each of the maps or files for these nodes is preprocessed according to which cell they are associated with as explained in Chapter 6.

Bayesian network

The function of the BN is to take all these state assignments and, using Bayesian statistics, answer certain operational questions. In this case, the probability of a rhino poaching for a specific area and time period is of interest. Other questions that can be answered are, what is the probability of a poacher being present, or what is the probability of a rhino being present given specified time, area, and context.

7.5 CONCLUSION

This chapter discusses the challenge of developing a BN, especially if it is attempted without the necessary software. Microsoft Excel is a well-known software tool, whereas Matlab is expensive and only used by researchers. A template tool was created in Excel to simplify the expert elicitation process whereby a BN structure can be developed quickly. It was especially effective in the workshop environment. Furthermore, the BN structure can be developed and simply emailed to other individuals so that they can populate the tables themselves.

The structure, the inner workings, and the software implementation of the latest model have been discussed in this chapter. In the next chapter the results and validation of the model will be discussed.

CHAPTER 8 RESULTS, VALIDATION, AND INSIGHTS

8.1 INTRODUCTION

This chapter evaluates how well the current expert knowledge model corresponds to the outcome of the expert workshop and to the belief of the experts. New insights gained during the testing phase are discussed, as well as their influence on the model. The model is altered slightly to incorporate these new insights, and what-if analyses are performed to test different scenarios.

8.2 NEW INSIGHTS TO THE MODEL

The conditional probability tables (CPTs) that were calculated in Chapter 6 are used to create a heat map of the situation, in particular, of the probability of a poaching event (*Poaching_event* = "True"). The changes that were made to the *Water*, *Distance_to_incursion_points*, and *Proximity_to_static_deterrents* variables when switching from 5×5 kilometres grid cells to 1×1 kilometre grid cells were also applied. The heat map that resulted from these CPTs did not completely correspond to the belief of the experts.

According to experts, as well as the rhino census data, rhinos are more prevalent in the centre and the south of the park. Analysing the rhino sighting data as well as the rhino poaching data, one would thus expect the likelihood of a poaching event to be higher in the southern and central parts of the park than in the northern parts. According to this map, however, the likelihood of a poaching event in the north is much higher than in the centre or the south. Owing to the sensitivity of the matter, the map can unfortunately not be displayed, but will be discussed nevertheless. A very light northern area coupled with light grey circles are due to the potential presence of poachers, as can be seen in Figure 8.1 (a). This is due to the proximity of the cells to static deterrents, as well as the distance to incursion points. The vegetation density is also weighted very strongly (Figure 8.1 (b)), which explains the black area on the northeastern side of the park.

There are also numerous dark spots on the map that coincide with the sighting of rhinos in past years (Figure 8.1 (c)). The weighting of the historical sightings of rhinos here is such that it weights a single sighting as strongly as it would weight, say, seven sightings occurring at the same location. Surely, if one rhino was seen several years ago, it should not carry as large a weighting as when several rhinos were seen at the same locations on different times. The landscape preference also seems to weigh the results too heavily, as it creates a medium-grey canvas on the central and southern parts of the park (Figure 8.1 (d)).

8.2.1 Changes to the model

The processing of *Historical_rhino_presence* was altered to more accurately reflect the clustering of the rhino sightings. A cell's eight neighbours were calculated and the total number of rhino sightings for those nine cells were counted. This number serves as the total number of sightings for that cluster. The distribution of the number of sightings per cluster was then used to compute the four states for classifying the cells of the map. Clusters where there were many rhino sighted now had a higher likelihood of containing a rhino in future, and cells where only a single rhino had been sighted had a much lower likelihood.

The anomalies in the heat map are attributed to incorrect weightings. These weightings were re-evaluated by reading the transcription notes of the expert workshop and thus going "back to the drawing board". Over the course of the project small details were overlooked or misinterpreted, and only by evaluating the heat map critically were these details discovered. The biggest changes occurred in the main nodes in each subgroup, namely the *Poacher_present*, *Rhino_present*, and *Poaching_event* nodes.

No significant changes occurred in *Active_deterrents* as there are currently no spatial variables included in that subgroup. The fact that there are no spatial variables in the *Active_deterrents* subgroup is also

(a) Potential presence of poachers

(b) Strongly weighted vegetation density

(c) Rhino sightings

(d) Strongly weighted landscape preference

Figure 8.1. Snapshots of the original poaching heat map

a point of contention. Spatial variables aim to focus the node on a specific location, as the BN is executed for each cell in the map. There should be an extra node in the *Active_deterrents* subgroup that links the cell location to the active deterrents. A possible solution to this problem would be to add a *Ranger_present* variable that states whether or not a ranger is present in that cell, but there are currently no data for this, making this solution impractical.

8.2.2 Changes to the poacher subgroup

Evaluating the heat map illuminated the fact that the *Vegetation_density* node seemed to have an incorrect weighting. Upon inspection it was found that the highest weighting had been given to *Vegetation_density* = "Open", as it was postulated that poachers will prefer the open areas and stay out of the very dense areas because there are no rhinos present in dense areas. Referring to the expert workshop notes and verifying the result with the poaching record data (as seen in Figure 6.11), it seemed that poachers indeed favour *Vegetation_density* = "Impenetrable vegetation" for poaching, but not for moving around. The simplified CPTs for the case when *Accessibility* = "High" and *Vegetation_density* = {"Impenetrable vegetation", "Open"} are presented in Table 8.1 and Table 8.2 respectively.

Table 8.1. Simplified CPT for *Accessibility* = "High", *Vegetation_density* = "Impenetrable vegetation"

Proximity_to static_deterrents		*Distance_to_incursion_points* 0-2.5 km	2.5-5 km	>5 km
		0.8000	*1.0000*	*0.2000*
0 - 5 km	*0.1000*	0.0784	0.0980	0.0196
5 - 10 km	*0.7000*	0.5488	0.6860	0.1372
> 10 km	*1.0000*	0.7840	**0.9800**	0.1960
Vegetation_density.Impenetrable vegetation				

Table 8.2. Simplified CPT for *Accessibility* = "High", *Vegetation_density* = "Open"

Proximity_to static_deterrents	*Distance_to_incursion_points* 0-2.5 km	2.5-5 km	>5 km
0 - 5 km	0.0470	0.0588	0.0118
5 - 10 km	0.3293	0.4116	0.0823
> 10 km	0.4704	0.5880	0.1176
Vegetation_density.Open			

The highest weighting should be given to *Vegetation_density* = "Impenetrable vegetation", as this is the scenario where poachers are less likely to get caught. The "best case" was given a baseline of

0.9800 instead of the previous 0.9000. The weighting of *Vegetation_density* = "Open" with respect to *Vegetation_density* = "Open" is 0.6000 as it is still not known whether poachers solely prefer extremely dense areas.

For *Vegetation_density* = "Impenetrable vegetation" the "ideal" state for the accessibility to be high was changed from *Distance_to_incursion_points* = "In cell" and *Proximity_to_static_deterrents* = "Far" to *Distance_to_incursion_points* = "2.5 - 5 km" and *Proximity_to_static_deterrents* = "> 10 km". The reason for this is that a poacher will try to stay close to incursion points, but if he is followed he might not use his intended incursion point, but rather one that is closer to his current location. The weightings between the states of *Proximity_to_static_deterrents* were changed from 0.1000, 0.5000, and 1 to 0.1000, 0.7000, and 1. The reason for this is that poachers will not poach close to a camp for fear of being noticed, they will try to stay as far away as possible.

The baseline for a poacher to be present when the accessibility is high, the weather conditions are rainy, the moon illumination is at its highest, and it is twilight, was adjusted upwards from 0.9000 to 0.9800 as this was deemed the best scenario for *Poacher_present* = "True". The reader is referred to Section 6.3 for a detailed discussion of the CPTs in the ranger subgroup.

8.2.3 Changes to the rhino subgroup

The baseline for rhino presence was adjusted from 0.9500 to 0.9800 when the time is preferred, the cell is close to water, and the number of sightings of rhinos in previous years is high. When *Landscape_preference* = "Preferred", the weightings between the states of *Water* changed from 0.3000 and 1, to 0.4000 and 1. The weightings between the states of *Historical_rhino_presence* changed from 0.0500, 0.2500, 0.7000, and 1 to 0.1000, 0.3000, 0.8000, and 1. The reason for the large difference between weightings is because the *Historical_rhino_presence* variable is now processed differently. The largest weighting now lies with the highest number of rhino clustering.

The weightings between the states of *Landscape_preference* were also changed slightly. For *Landscape_preference* = "Neutral" the weighting was changed from 0.7000 to 0.6000 to give slightly less weight to neutral areas, although neutral areas are used when there are no preferred areas nearby.

8.2.4 The new heat map

The latest expert knowledge heat map of poaching events can unfortunately not be shown due to the sensitivity of the problem, hence it will only be described in the text. The overall colour of the map is very dark, signifying that a poaching event is rare. The northwestern part of the park is slightly lighter than the north eastern part as can be seen in Figure 8.2(a). This is due to the denser vegetation in the east and also fewer poaching events. There are also more incursion points and less static deterrents in the west which adds to the lighter shade and thus the higher probability of a poaching event. The central and southern parts of the park are speckled with white and light greys patches, corresponding to areas of high poaching event predictions (Figure 8.2(b)). The lighter areas do not extend to the southern borders, as there are numerous static deterrents on the southern borders.

(a) Denser vegetation in the east of the park

(b) High poaching event predictions

Figure 8.2. Snapshots of the original poaching heat map

The heat map is broken down into a map for *Poacher_present* (Figure 8.3) and a map for *Rhino_present* (Figure 8.4). Analysing Figure 8.3 it appears that the likelihood for a poacher being present is more or less the same in the northern part of the park as in the southern part. In this map *Poacher_present* is a culmination of the spatial variables *Vegetation_density*, *Distance_to_incursion_points*, and *Proximity_to_static_deterrents*. The circular patches are due to the fact that both *Distance_to_incursion_points* and *Proximity_to_static_deterrents* are processed by calculating the distance between a cell's midpoint and the closest incursion point or static deterrent. The borders on the eastern side, as well as the northwestern and southwestern side, could be preferred by poachers and can be seen as weak spots on the border. The areas of high poacher presence coincide

Figure 8.3. *Poacher_present* prior map

Figure 8.4. *Rhino_present* prior map

with known weak borders: the park shares its entire eastern border with Mozambique, the northernmost part of the park is close to the Zimbabwean border, and the southwestern border is close to informal settlements.

A poacher will keep away from populated places and fixed structures, but will prefer to be close to incursion points, thereby entering and exiting the park more easily. The slightly darker areas in the south are due to more static deterrents in a smaller area and the circles overlap. Most of the rhino population is clustered in the southern part of the park, thus there would be an increase in security.

The *Rhino_present* prior map is the culmination of the spatial variables *Landscape_preference*, *Water*, and *Historical_rhino_presence*. The white patches on the map are mainly due to the influence of the historical rhino sighting data. The vegetation of the KNP lends itself to the slight partitioning seen in Figure 8.4 between east and west in the northern part of the park, and corresponds to the rhinos' preference in terms of grazing and wallowing. Rhino sightings are less clustered in the northern part of the park when compared to the southern part. The dark area going from north to south on the eastern side of the KNP, as well as the two small dark patches in the south, are due to the rhinos' avoidance of those landscapes, either due to a scarcity of food, wallowing pans, or too many areas frequented by humans.

Heat maps can be drawn for each spatial variable represented in the poaching event heat map. The maps for *Distance_to_incursion_points* and *Proximity_to_static_deterrents* are not shown due to the sensitivity of both the locations of the deterrents as well as the incursion points. The *Vegetation_density* and *Landscape_preference* maps are presented in Figure 8.5 and Figure 8.6 respectively. The black shading in Figure 8.5 corresponds to areas with dense vegetation, and the white shading corresponds to open areas. The black shading in Figure 8.6 corresponds to avoided areas, the grey shading corresponds to neutral areas, and the white shading corresponds to preferred areas. The reader is referred to Section 6.3 and Section 6.4 for in-depth discussions of these variables.

The *Water* prior map is shown in Figure 8.7 and Figure 8.8. The isolated black patches in Figure 8.7 are the areas in the KNP that are further than five kilometres away from a water source. Figure 8.8 shows a zoomed in image of the water map together with the water sources and main rivers to show the detail.

8.3 EVALUATING THE RESULTS

The model now corresponds to the experts' belief as far as the heat map is concerned. It is difficult to accurately evaluate the model when there is no right or wrong answer in terms of the outcome of the model. One way of checking the validity of the model is to perform "what-if" analyses or queries to test whether or not the likelihoods correspond to the experts' belief. In this section six of these queries are performed by the author and discussed. The concept of counterfactuals as described in Chapter 3 is also revisited.

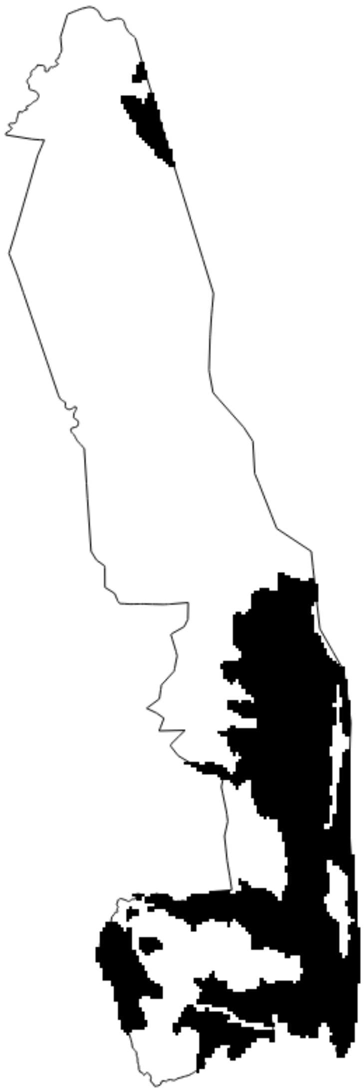

Figure 8.5. *Vegetation_density*

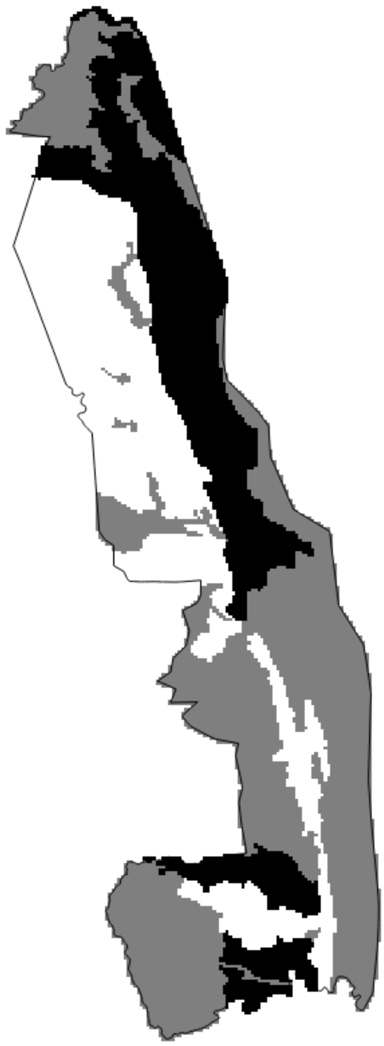

Figure 8.6. *Landscape_preference*

As mentioned earlier, the model presented in this thesis is an expert-driven model, and was developed because high-quality and complete datasets were not available at the time. Since then, the focus on data gathering has intensified and the quality and quantity of data is now much better than it was at the start of this study. However, owing to the fact that this model is a purely expert-driven model, it cannot easily be trained on the available data. Experts tend to add abstract variables to BNs in order to simplify it, but the challenge is usually that there does not exist data for these abstract variables (think about *Accessibility*, for instance). The challenge then is, how do you test your model that is parameterised on expert knowledge, and not on data? According to Sargent [108] there are many ways to validate a model, and as Carley [107] puts it, "Part of the argument, as will be seen, is that not all

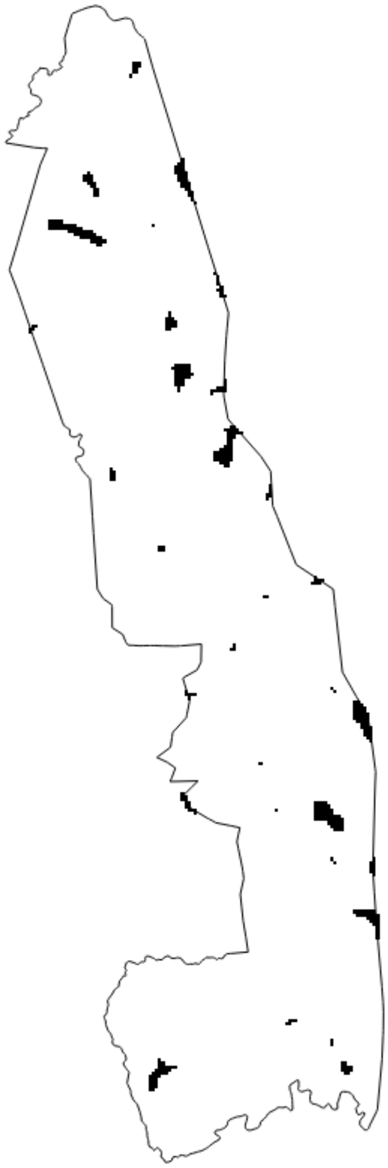

Figure 8.7. *Water* prior map

models need to be validated and that the level of validation chosen depends on the model's purpose." One validation approach is to use case-based evaluation to verify the validity of the expert-driven model [106].

8.3.1 Query 1: Predictive mode

Chapter 3 mentioned that Bayesian inference can occur in two directions: top-down (predictive) and bottom-up (diagnostic or prescriptive) [37]. In this study, *Poaching_event* is the target variable and is

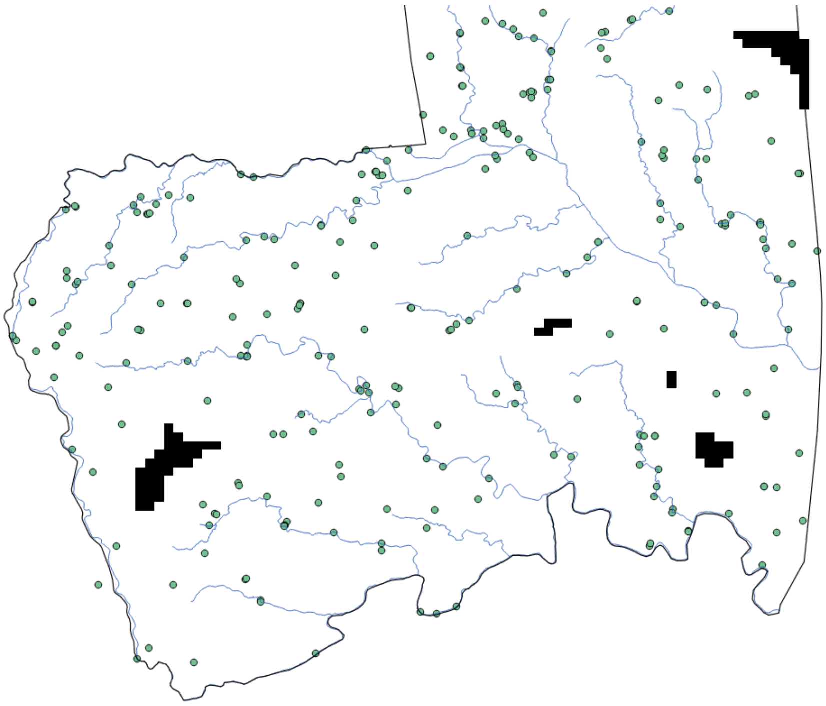

Figure 8.8. *Water* map zoomed in

at the bottom of the network. The goal is to predict rhino poaching events, thus a top-down or predictive approach is followed. The discussion is started with best- and worst-case scenarios for poaching events to occur. In each case the states of the variables are deliberately chosen to showcase the absolute worst (first example) and absolute best (second example) outcome possible for the model.

Worst-case (Least Advantageous) Scenario

For the worst, least advantageous case we would like to know which scenario of variable states would maximise the probability of a poaching event. This optimal set of states is given in the first two columns of Table 8.3 and the Hugin® inference screen is shown in Figure 8.9.

The probability of a poaching event occurring is now 87.60%. Setting the root nodes to their optimal states for a poaching event gives a high probability for *Poacher_present* = “True” (96.63%), a high probability for *Rhino_present* = “True” (98.00%), and a high probability for *Active_deterrents* = “Not effective” (88.80%) which is in line with RAT. The probability of a poacher being present is dependent on a high probability for *Accessibility*. The optimal factors for creating a high probability

Table 8.3. Worst-case and Best-case

Variable	State (worst case)	State (best case)
$Vegetation\_density$	Impenetrable vegetation	Open areas
$Distance\_to\_incursion\_points$	2.5-5 km	> 5 km
$Proximity\_to\_static\_deterrents$	> 10 km	≤ 5 km
$Time\_of\_day$	Twilight	Midday
$Moon\_illumination$	75 - 100%	0 - 25%
$Weather$	Rainy	Clear
$Landscape\_preference$	Prefer	Avoid
$Water$	Close	Far
$Historical\_rhino\_presence$	$[39, \infty)$	$[0, 13)$
$Corruption\_index$	Corrupt	Not corrupt
$Stages\_of\_the\_month$	Payday + 7	Rest of the month
$Festive\_periods$	True	False

of $Accessibility$ = "High" is when the vegetation is dense, the cell is between two and a half and five kilometres from an incursion point, and when the cell is more than 10 kilometres from a static deterrent. The rest of the optimal states for a poacher to be present are when it is twilight with between 75% and 100% moon illumination, and when it is rainy. When all these factors are in place, the probability of $Positive\_poacher\_observation$ = "True" changes from 8.09% to 38.82%. The reason for this is that the presence of a poacher increases the probability of observing that poacher.

The optimal scenario for a rhino to be present is when the cell contains landscape that is preferred by the rhinos, the cell is within five kilometres from water, and when there have been a total of at least 39 rhino sightings in that cell during the past four years during the annual census count. These factors are in line with our belief that rhinos prefer to stick to certain areas, and that they prefer to be close to water, preferred grazing areas, and wallowing pools. All these factors lead to $Rhino\_track$ = "True" changing from 7.81% to 78.40% and $Rhino\_sighted$ = "True" changing from 10.96 to 19.80%. The reason for this is that the presence of a rhino increases the probability of obtaining a rhino track, as well as increasing the probability of seeing that rhino. The ideal situation for the active deterrents to be ineffective is during a festive period, when it is the week after payday, and when the ranger is corrupt.

When the probability of *Poaching_event* = "True" seems low but the conditions for a poaching event seem favourable, it has to be remembered that each probability is for a single cell. Cell i might have a likelihood of 10.96% for a poaching event, but cell $i+1$ might have a likelihood of 54.09%. Cell $i+1$ is thus much more likely to have a poaching event than cell i. A rhino poaching event is also viewed as a rare event and is not an event that has a high prior probability of occurring.

Best-case (Most Advantageous) Scenario

The best case refers to the combination of circumstances for a rhino poaching event to be the least probable. We are thus not setting the target variable *Poaching_event* = "False", we are simply looking at the factors individually to decide how to stop a poaching event. Table 8.3 illustrates the variable settings to obtain this in comparison to the case where a poaching event is definitely going to occur as seen in the previous explanation. Figure 8.10 presents the Hugin® inference screen.

The probability of a poaching event is now zero (0.0000707%). The probability of *Poacher_present* = "True" is extremely low (0.0300%), as is the probability of *Rhino_present* = "True" (0.3900%). The probability of *Active_deterrents* = "Not effective" is also quite low (20.00%), which aligns with RAT. The optimal state for a poaching event to not occur is the opposite of what it is for a poaching event to occur. Therefore, if a variable has two states and state 1 was part of the optimal scenario for a poaching event to happen, then state 2 will be part of the optimal scenario for a poaching event to not happen. Examples of this are *Vegetation_density*, *Water*, *Corruption_index*, and *Festive_periods*. Examples of where the optimal answer was the opposite extreme state of the variable are *Proximity_to_static_deterrents*, *Moon_illumination*, *Weather*, *Landscape_preference*, and *Historical_rhino_presence*. In the worst case the following variables will be: *Distance_to_incursion_points* = "≤ 5 km", *Time_of_day* = "Midday", and *Stages_of_the_month* = "Rest of the month".

8.3.2 Query 2: Diagnostic mode

In a diagnostic approach the conditions are inferred that need to be in place to ensure that the target variable occurs. The question is thus, what should the states of the variables be so that a poaching event will occur? For this inference mode the target variable's value is fixed: the node *Poaching_event* is set to "True". Figure 8.11 illustrates the Hugin® inference screen.

If *Poaching_event* = "True" we can see that both *Poacher_present* and *Rhino_present* have the highest likelihood of being true, since there cannot be a poaching event without a poacher and a rhino. The likelihood for *Active_deterrents* = "Not effective" is 66.02% which is in line with Routine Activity Theory (RAT) that states the three elements for a crime to occur. For a poacher to be present, the accessibility into the park needs to be high. According to these numbers, the *Accessibility* can be low, but the vegetation density must be open, the poacher must be far from incursion points (the rangers know where the incursion points are, thus it will be better to steer clear of them for as long as possible), and at least five kilometres from static deterrents such as camps and ranger posts. Twilight and night are also the preferred times of day, as is a high moon illumination percentage which maximises their ability to see. The combination of clear and rainy conditions seems to be preferred, which contradicts the notion that poachers prefer to poach during either windy or rainy conditions. A contention is that clear weather gives the poachers better vision which is beneficial in negotiating the terrain, avoiding the rangers, and making a hasty escape.

For a rhino to be present the landscape has to be at least neutral, meaning that the rhinos will graze there if there is nothing else to eat nearby. They also have to be close to water. The historical presence node indicates an interesting outcome, with the rhinos not having to be highly clustered in a cell. The reason for this could be that it does not matter to poachers if there is a single rhino or five rhinos. What matters most is that there is at least one rhino present. Another reason for this is that the poachers possibly avoid areas with a large concentration of rhinos since they fear that that area might be better patrolled than areas with only one or two rhinos.

For the active deterrents to be ineffective it seems that it should not fall over a festive period, which is an interesting outcome. This could either mean that the experts were erroneous in their assumption that festive periods play a role, or that the festive period states should have been calculated differently. The proximity to payday also does not seem to make any difference as the probabilities for each state is more or less a quarter. The "best case" is also when the ranger is not corrupt, which also goes against the belief. A possible explanation for this is that, as already mentioned, the *Active_deterrents* subgroup does not contain any spatial variables, only temporal variables. The *Active_deterrents* subgroup needs to be revised, as the experts could possibly have been incorrect in their choice of important variables, and the influence of those variables on the rangers' effectiveness. There could be other variables influencing the effectiveness of active deterrents. Conversely, the active deterrents do not make a difference at all.

8.3.3 Query 3: Prescriptive mode/counterfactual

In a bottom-up or prescriptive approach the conditions are inferred that are needed to prevent the target variable from occurring, given the context. In other words, what should the states of the variables be so that a poaching event does not take place? The leaf nodes are observed and the causes are inferred, as per the explanation of the counterfactual in Chapter 3. This mode is almost the opposite of the diagnostic mode: we want to know how to prevent a poaching attack, thus *Poaching_event* = "False". Figure 8.12 illustrates the Hugin® inference screen.

The case that is now shown is how to stop a poaching event, according to the experts. The likelihood of *Poacher_present* = "True" is now very low (8.25%), as the best way to prevent a poaching event is to prevent a poacher from being present. The likelihood of rhinos being present is also very low (9.08%), as a rhino needs to be present for a poaching event to occur. The likelihood of active deterrents being effective is the same as for active deterrents to not be effective, thus it does not seem to make a difference in stopping poaching events. An explanation for this could be that rangers and other active deterrents are always present in the park, and that a change in their effectiveness might not be deliberate or due to corruption. It could also mean that the effectiveness of active deterrents do not make any difference in whether or not a poaching takes place. The absence of any spatial variables in the subgroup could also make a difference in terms of the influence it has on the occurrence of a poaching event.

8.3.4 Query 4: Impossible case

For a rhino poaching event to occur both a rhino and a poacher need to be present. Therefore there are two cases when a poaching event is impossible: either there is (1) no rhino at that location, or there is (2) no poacher there. For these situations, we would expect the probability of a poaching event to be zero. Figure 8.13 illustrates the first case, and Figure 8.14 illustrates the second case. For both cases it can be seen that the probability of a poaching event is zero.

8.3.5 Query 5: Predictive mode (another example)

The last example shows the case where the probability is required of a poaching event occurring ($Poaching\_event$ = "True") given certain information. Below is a scenario that might arise if the tool is used in the park.

Imagine the following scenario in the KNP: The commander receives intel that there will be a poaching event somewhere in the KNP the next day. He received certain pieces of information from a source and he wants to use the predictive model to check if the information is reliable. The weather for the next day is predicted as overcast and the moon illumination percentage will be 40%. It is far from payday and it does not fall within a festive period. He knows that the poachers will strike in an open area that is not too close, but also not too far away from exit points in the fence. He also knows that the poachers will steer clear of camps and ranger patrol huts. The poachers will strike around last light close to water where many rhinos have been seen in recent years, yet the location is not necessarily preferred by rhinos in terms of grazing. He also knows that the ranger responsible for that section of the KNP might have a connection to a poaching syndicate. The commander can now populate the model with the information above, given in table form in Table 8.4.

Table 8.4. What-if analysis: Predictive

Variable	State
$Vegetation\_density$	Open areas
$Distance\_to\_incursion\_points$	2.5 - 5 km
$Proximity\_to\_static\_deterrents$	5 - 10 km
$Time\_of\_day$	Twilight
$Moon\_illumination$	25 - 50%
$Weather$	Overcast
$Landscape\_preference$	Neutral
$Water$	Close
$Historical\_rhino\_presence$	$[13,26)$
$Corruption\_index$	Corrupt
$Stages\_of\_the\_month$	Payday + 14
$Festive\_periods$	False

The probability for *Poaching_event* = "True" is 3.48%. Figure 8.15 illustrates the Hugin® inference screen. The probability of good accessibility into the park is 41.16%. According to the experts, the poachers prefer dense vegetation and steering clear of built-up areas such as camps. In this case *Vegetation_density* = "Open areas" and *Proximity_to_static_deterrents* = "5 - 10 km", thus decreasing the probability of *Accessibility* = "High". A lower probability of good accessibility also greatly influences the probability of whether or not a poacher is present. Here the probability of *Poacher_present* = "True" is 23.05%, which is quite low. According to the experts, poachers prefer to poach during a high moon illumination and during rainy weather conditions. In this case *Moon_illumination_percentage* = "25 - 50%" and *Weather* = "Overcast", thereby decreasing the probability of *Poacher_present* further.

The probability of *Rhino_present* = "True" is 17.64%, which is also quite low. We would expect that there would be more rhinos in their preferred areas than in the neutral areas, where they will only graze if there is no other food available. According to the experts, rhinos tend to stay in certain areas, so if they have been seen at location X in the past three years, they will most probably be there this year too. Thus we would expect that rhinos would be seen more in the areas of high clustering. In this case *Landscape_preference* = "Neutral" and *Historical_rhino_presence* = "$[13,26)$", thereby reducing the probability of *Rhino_present*.

The probability of *Active_deterrents* = "Effective" is 25.60%. The three contributing factors to whether or not the active deterrents are effective are, according to the experts, whether or not it is a festive period, whether or not the ranger is corrupt, and what stage of the month we are in. A corrupt ranger drastically decreases the probability of the active deterrent's effectiveness. According to the experts, the active deterrents are more likely to be ineffective over festive periods and close to payday, thus the fact that *Festive_period* = "False" and *Stages_of_the_month* = "Payday + 14" decrease the probability that *Accessibility* = "Not effective" slightly.

According to RAT a poaching event can take place when there is a poacher present (P(*Poacher_present* = "True") = 0.2305), a rhino present (P(*Rhino_present* = "True") = 0.1764), and an absent guardian (P(*Active_deterrents* = "Not effective") = 0.7440). There is thus a chance that a poaching event can take place, but it is very small (P(*Poaching_event* = "True") = 0.0348).

The commander can now decide if he agrees that this is a high probability for a poaching event (remember that a poaching event is a rare event and that the probability for a cell is not absolute, it should be viewed in relation to the cells next to it). If, say, the cells surrounding that cell have a probability of 1.02% the commander can say that the cell he is studying seems to be three times as likely to contain a poaching event.

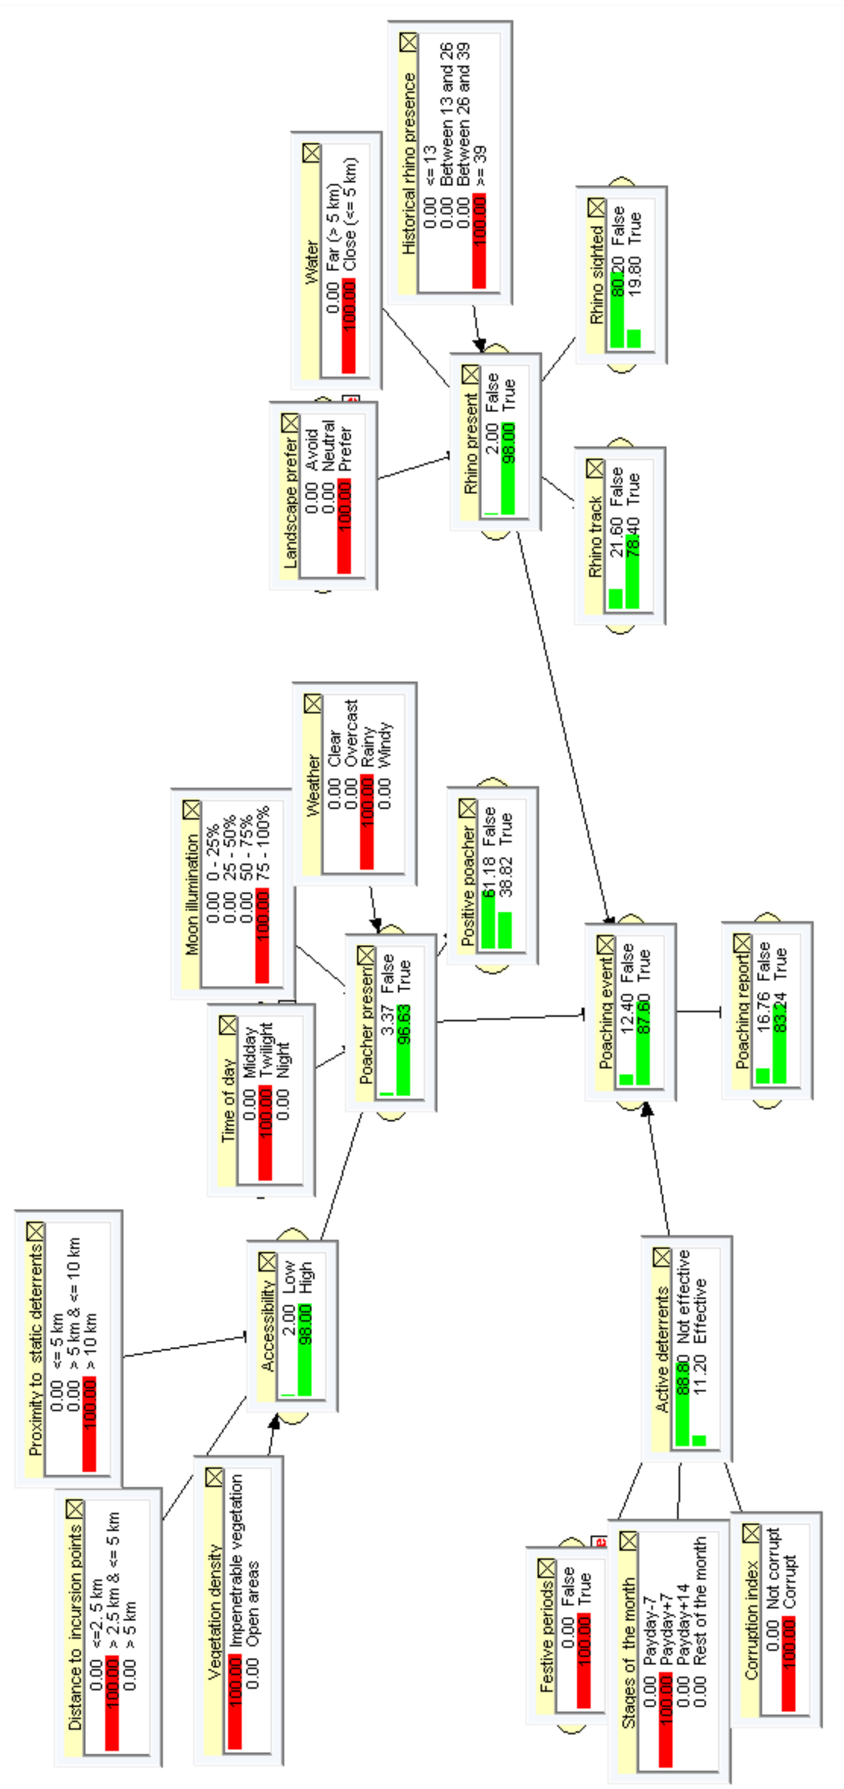

Figure 8.9. What-if analysis: Worst-case

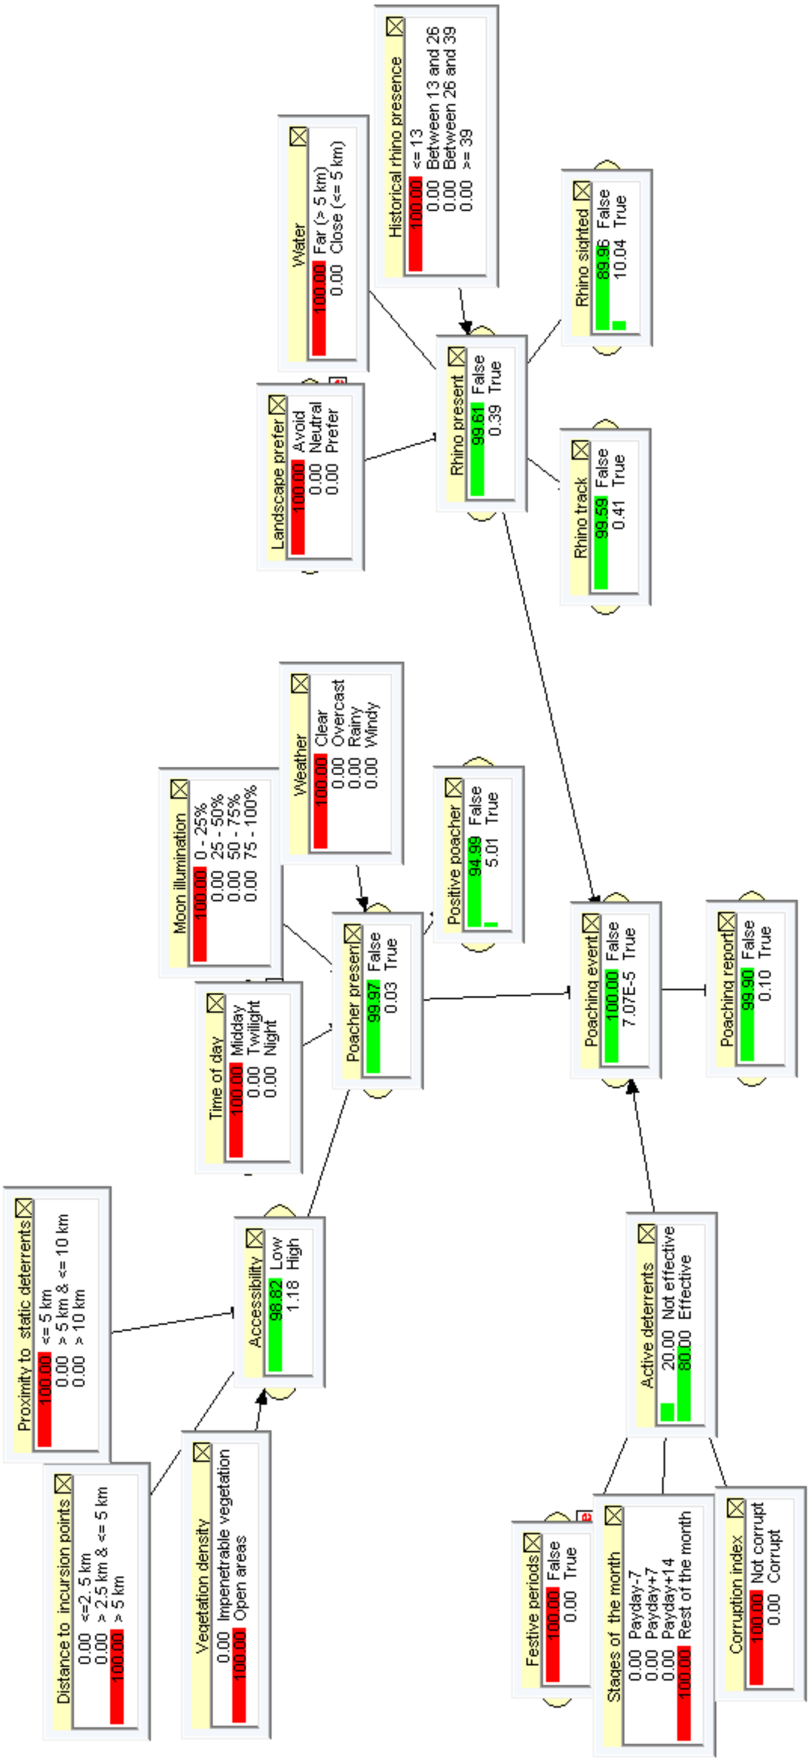

Figure 8.10. What-if analysis: Best-case

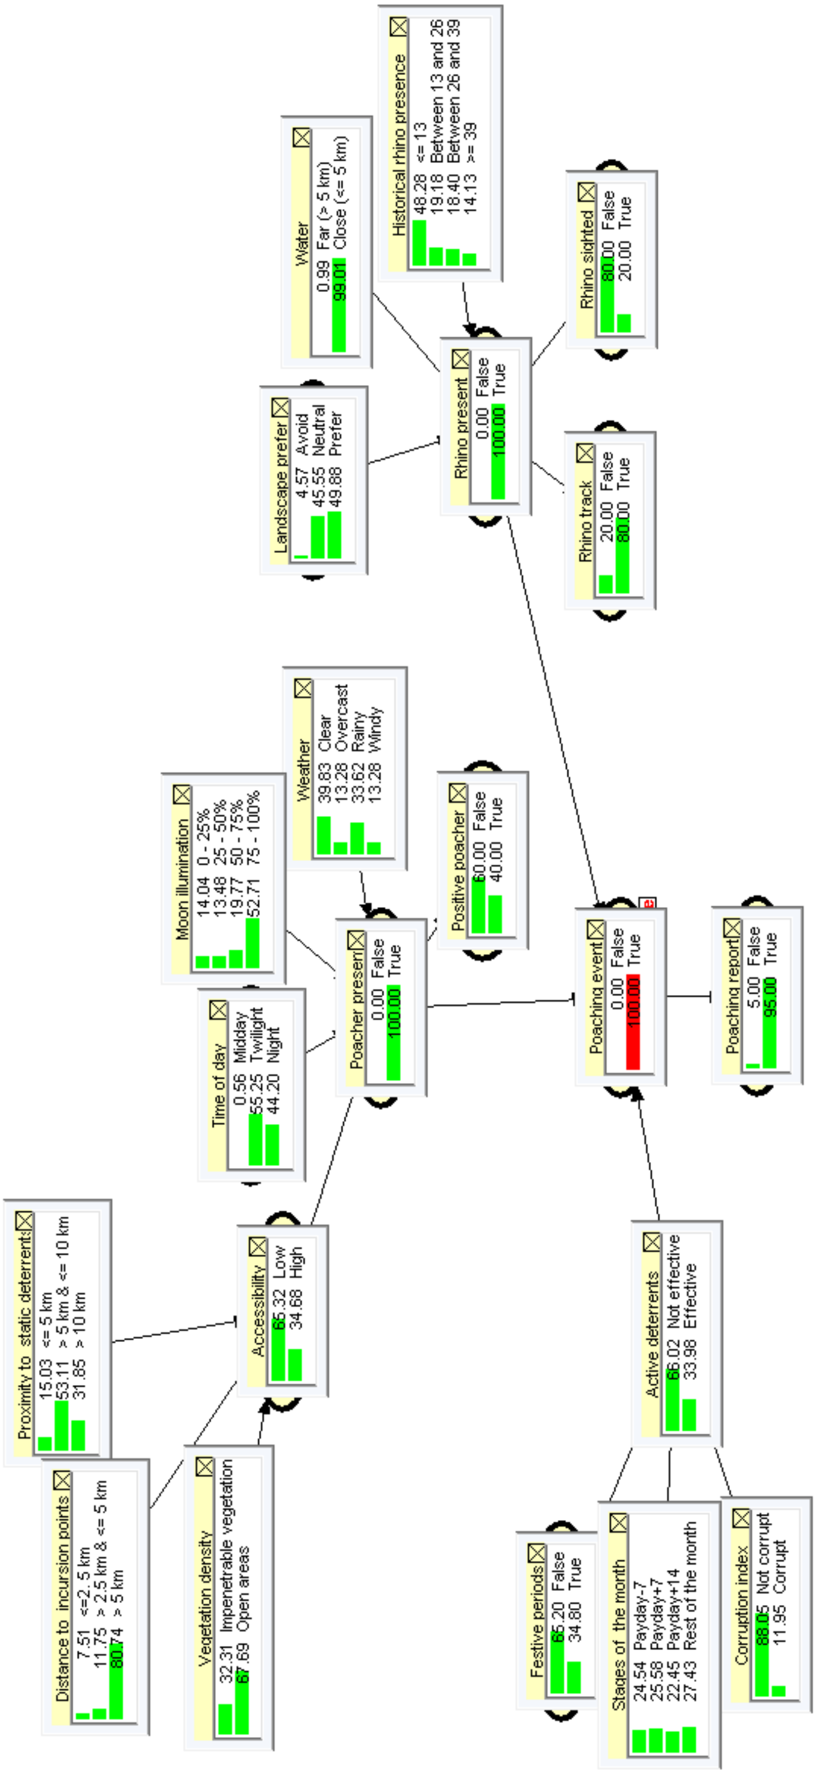

Figure 8.11. What-if analysis: Diagnostic

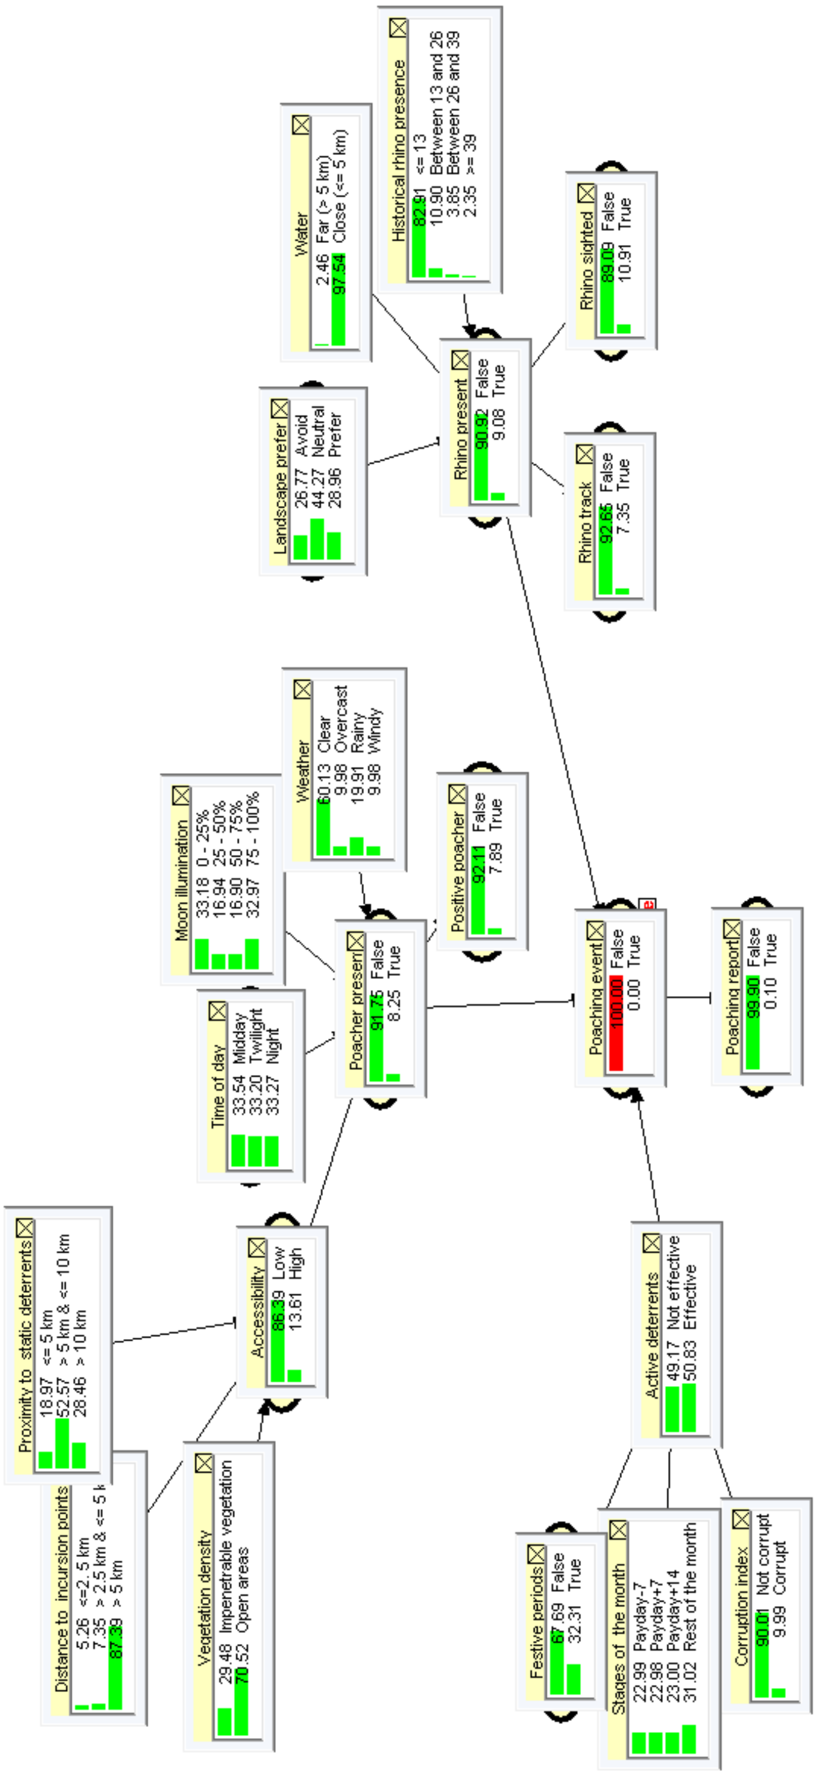

Figure 8.12. What-if analysis: Prescriptive

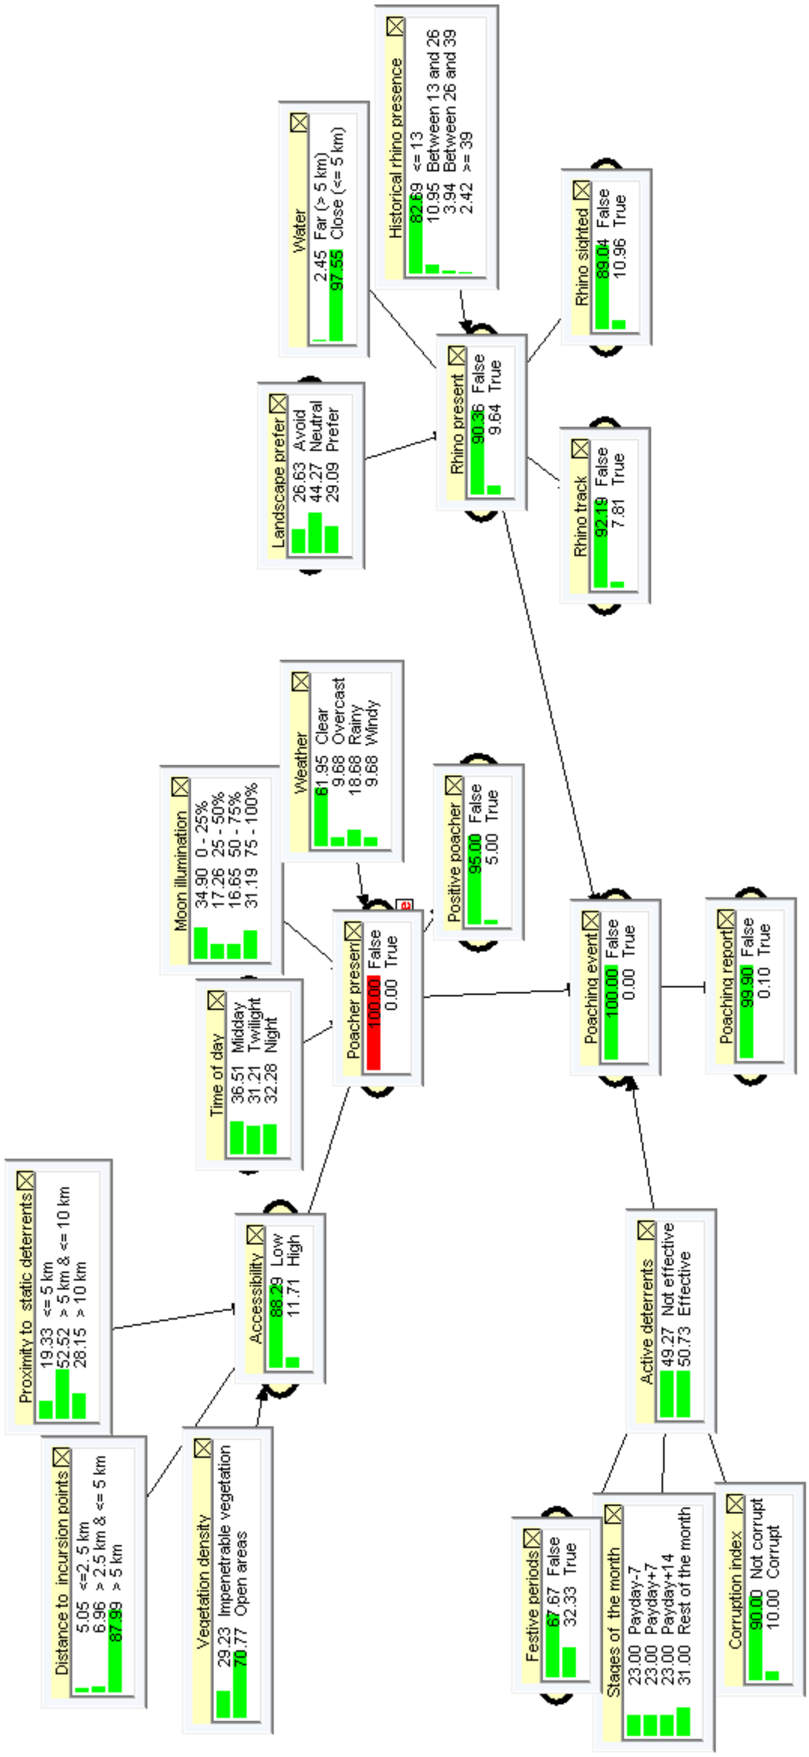

Figure 8.13. What-if analysis: Impossible (1)

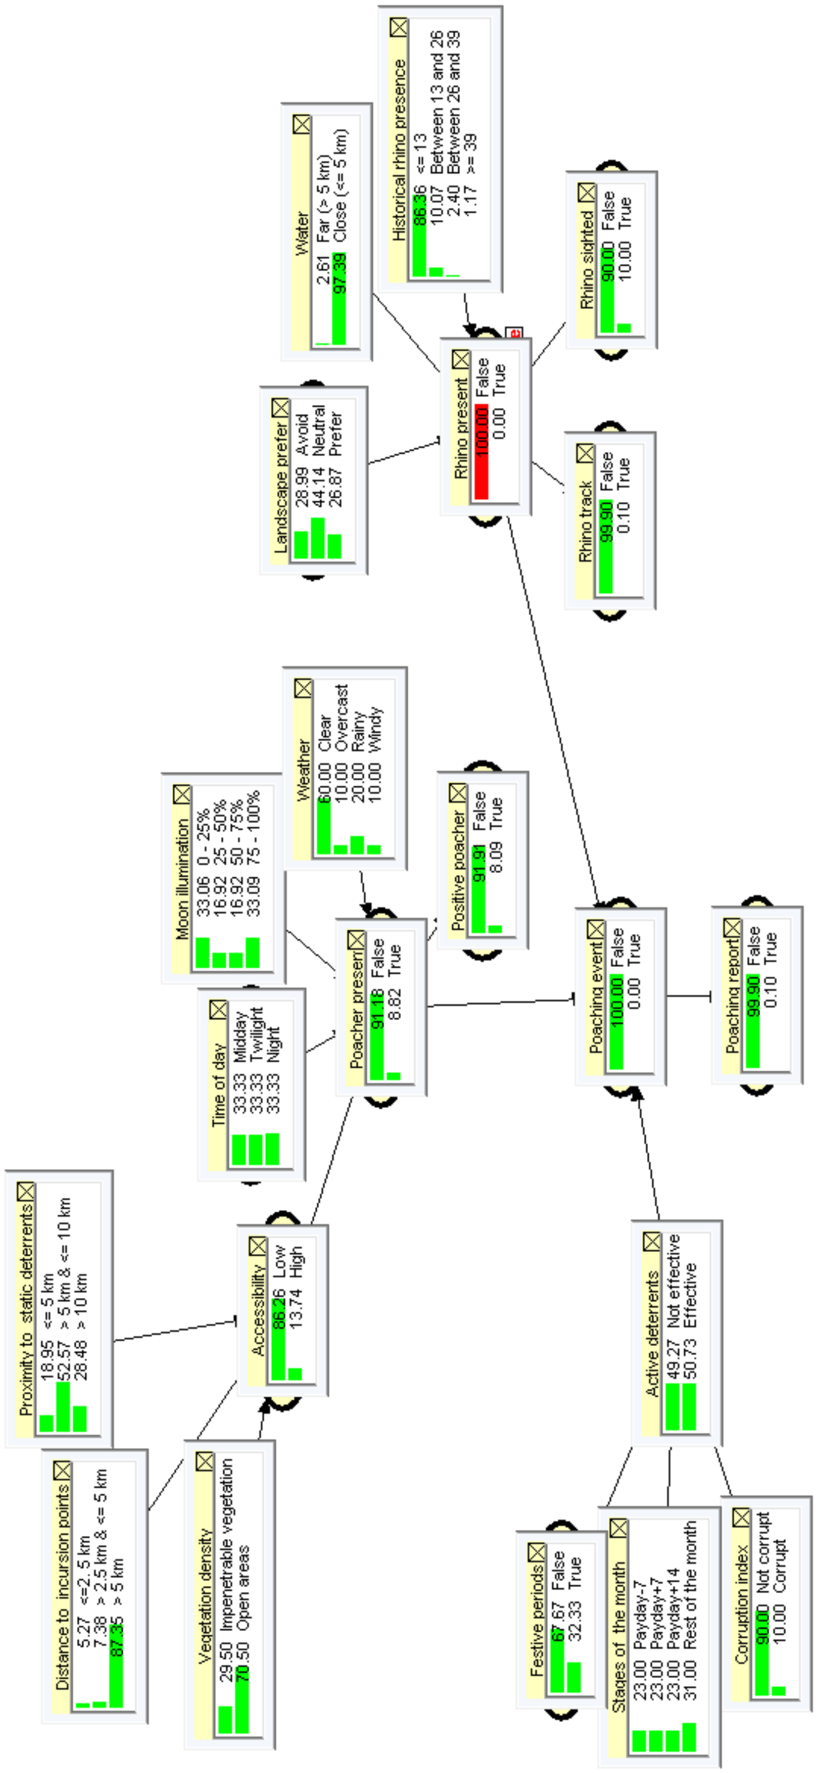

Figure 8.14. What-if analysis: Impossible (2)

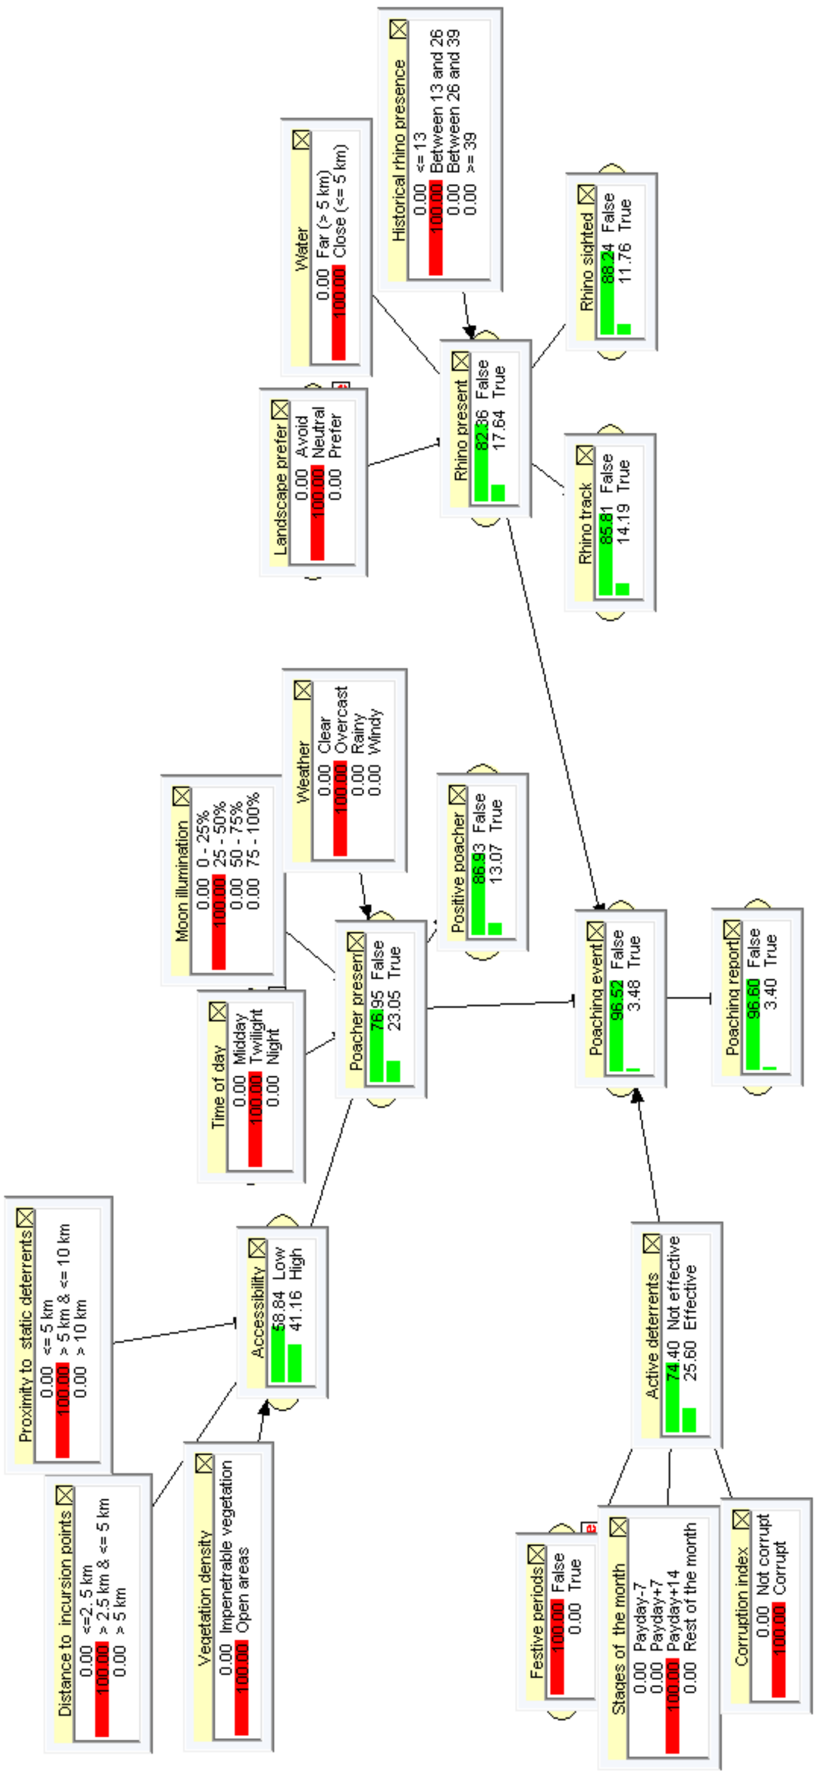

Figure 8.15. What-if analysis: Arbitrary case

8.4 INDEPENDENT EXPERT EVALUATION

The above what-if analyses were shown to experts in the KNP and they were asked to evaluate it in terms of the outcomes of the analyses as well as the model structure. 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. Mostly they agreed with the model architecture, but suggested a few adjustments. According to them, the model reflects the situational picture roughly a year ago. The variables remain the same, but the important states for a poaching event to occur has changed somewhat. Owing to the sensitivity of the problem all the details cannot be divulged, but some of the adjustments and improvements are described below.

The experts started evaluating 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 equal weight, but dense vegetation seemed to have been 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 times are associated with poaching events, thus it is important to distinguish between morning and evening.

The experts agree that the poachers' methods of attacking differ between north and south. One node influenced by this is the *Moon_illumination_percentage* where a low illumination percentage is preferred in one area of the park, and full moon is preferred in another area of the park. The experts suggested that the states for *Moon_illumination_percentage* rather be "Quarter moon", "Dark moon", and "5 days before full moon until 5 days after full moon", but it was explained to them that these states boil down to our current states for *Moon_illumination_percentage*.

The what-if analysis for the "Best case" contained the combination of *Moon_illumination_percentage* = "75 - 100%" and *Weather* = "Rainy", which, according to the experts, will not happen. The rainy or overcast conditions are only important during the day. Rainy conditions would cancel out the fact that the poachers are there during a high moon illumination percentage. The combination of *Moon* = "75 - 100%" and *Weather* = "Overcast" is assumed to be a low-probability combination. The sound of a shot travels roughly twice as far during overcast weather thus increasing the poacher's chances of being caught. 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.

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 prefer twilight and full moon, and the other might prefer the middle of the day and a dark moon at night. A node such as this might be useful in profiling the poachers, but using it in this model will require 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, 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. This is not necessarily corruption or malice, but more likely ill-discipline. It could also be that the ranger is incapable due to his experience level, the fact that he is on leave, or is exhausted.

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 what-if analyses. The rangers are dedicated and work long hours, even more so than usual during festive seasons, thus reducing the time for taking leave or spilling secrets in a social situation. A possible problem could be when leadership takes leave and is away from the park.

The experts concur that they lose communication with their staff over Christmas and Easter. They have to grant half of the staff leave over these festive periods, thus they are understaffed. Christmas of 2015 was one of the first Christmas periods in the past few years where the poaching total did not increase exponentially. The staff was optimised over Christmas from mid-November until mid-January and they managed to subdue the number of poaching attacks. According to the experts, festive periods are not really a problem, but if poaching events are not actively managed during those times, the months following that festive period will be hit harder than normal.

8.5 VALIDATION

Users of a model, and individuals influenced by a model's outcomes, have a right to question whether or not the model in use is correct and behaves as it should [108]. "Putting their minds at ease" happens by means of verification and validation. Verification verifies that the model is fit for its intended purpose, whereas validation ensures that the outputs are correct. In the previous sections validation was performed through *extreme condition testing* (Query 4: Impossible case), *face validity* (experts in the KNP were asked whether the model behaves as it should), and *parameter variability-sensitivity analysis* (the what-if analyses), which are all, according to Sargent [108], legitimate ways of validating a model.

Previously in the thesis it was explained why the model was kept free from data, and also why it is difficult to train and test the model on data. There does not currently exist a method to adequately test this type of model on data, but for the sake of interest a goodness-of-fit test is performed to illustrate this shortcoming.

8.5.1 Goodness-of-fit test

The chi-square goodness-of-fit test evaluates the validity of a distribution based on a hypothesis. The null hypothesis is evaluated against the alternative hypothesis. In this case we want to test whether or not the poaching data could have come from the expert-driven model.

The test statistic for the chi-squared test is given by:

$$X^2 = \sum_i \frac{(O_i - E_i)^2}{E_i}, \tag{8.1}$$

where O_i is the observed value in the ith section and E_i is the expected value in the ith section. The observed values refer to recorded poaching events and the expected value to predicted poaching events.

8.5.2 Validating the expert model

Validating the expert-driven model is problematic as the model contains many abstract nodes for which there are no data available (such as *Accessibility*, for example). The rule of thumb for a goodness-of-fit test is that the expected frequency (in this case, per cell) should be at least five in order for the approximation to be valid [109]. Many cells in the grid do not contain any data for poaching events, thus validating the map on a cell-by-cell basis will not work in its current configuration.

In order to use a goodness-of-fit test statistic, the map was divided into the 22 ranger sections of the KNP so that each section has a fair amount of poaching events. Figure 8.16 illustrates the sections each with its own share of the 3,000 poaching event data points. Owing to the sensitivity of the problem the poaching data cannot be shown, but the number of poaching data points per ranger section is shown on the map. This poaching data was collected over a few years and, again, owing to the sensitivity, the location as well as the yearly distribution can unfortunately not be divulged.

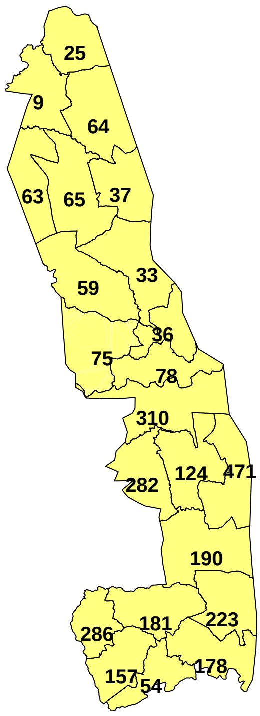

Figure 8.16. Poaching events per ranger section

The poaching event data points are the observed values and the predicted poaching events are the expected values. This needs to be calculated per ranger section, thus the total number of poaching event data points per ranger section was calculated. These poaching counts were divided by the area of each section to obtain a poaching count per section per square kilometres. The number of 5×5 kilometre cells overlapping with each of these sections was then calculated. If a cell overlaps with two sections, the cell is awarded to the section with which it overlaps the most. The normalised heat map of predicted probabilities was then used to compute the prediction value for a section by adding the probabilities of the cells overlapping that section. So, if Section 5 has a total of four cells overlapping with it and, according to the normalised heat map, the cell probabilities are 0.0041, 0.2100, 0.0071, and 0.0333, the prediction value for Section 5 will be 0.0041 + 0.2100 + 0.0071 + 0.0333 = 0.2545. Thus, all the cells in the map were allocated to a ranger section and their corresponding cell probabilities were summed for each of these ranger sections. All of these section probabilities were multiplied by 3,000 (the total number of poaching event data points) and divided by the area of the section in square kilometres in order to compare it to the observed poaching values. Table 8.5 shows the observed poaching values next to the expected poaching values (both sets of numbers are per section per square kilometre).

The test statistic is evaluated by testing whether the null hypothesis (H_0) should be accepted or rejected, based on a specific α-value. In this case H_0 states that the poaching data came from the expert-driven model and the alternative hypothesis (H_1) states that the poaching data did not come from the expert-driven model. The degrees of freedom are $(k-1) = 21$ degrees where $k = 22$ are the regions used in the calculations.

From Table 8.5 it can already be seen that the values look reasonably close to each other, thus it might be that they came from the same distribution. According to Wackerly [109], if α is chosen as $\alpha = 0.1$ the null hypothesis should be rejected when $X^2 > 29.6151$. Applying Equation 8.1 to the data in the first two columns of Table 8.5 yields $X^2 = 1.9306$. X^2 is thus much smaller than the tabulated critical value of χ^2, thus the null hypothesis is not rejected. The observed poaching values might thus have come from the expert model.

The same test can be performed for the uniform model where each cell in the map has an equiprobable chance of occurring. The uniform model is also a model, albeit a very bad one. Here the null hypothesis states that the observed poaching values could have resulted from the uniform model. The uniform values per section and per square kilometre can be seen in the third column of Table 8.5. The value for α is again chosen as $\alpha = 0.1$, thus the null hypothesis should be rejected when $X^2 > 29.6151$. The goodness-of-fit test yields $X^2 = 2.8481$, thus the null hypothesis is not rejected. The X^2-value is,

Table 8.5. Observed and expected poaching event values

Section	Observed	Expected	Uniform
1	0.0336	0.0284	0.1688
2	0.0465	0.0050	0.1407
3	0.0320	0.0230	0.1479
4	0.0645	0.0357	0.1641
5	0.2703	0.1824	0.1803
6	0.5159	0.1622	0.1559
7	0.3673	0.2911	0.1766
8	0.1312	0.2509	0.1363
9	0.1620	0.2845	0.1648
10	0.5461	0.2659	0.1748
11	0.2321	0.2481	0.1604
12	0.1017	0.1509	0.1916
13	0.2147	0.2341	0.1447
14	0.0782	0.1590	0.1682
15	0.0114	0.0346	0.1754
16	0.0564	0.1625	0.1295
17	0.0595	0.0215	0.1606
18	0.0505	0.1296	0.1460
19	0.0724	0.1617	0.1572
20	0.2810	0.2735	0.1495
21	0.2157	0.2566	0.1725
22	0.1121	0.1204	0.1412

however, higher than that of the expert model, thus the expert model was the more likely source of the two of the observed poaching values.

The visual results of the goodness-of-fit tests are illustrated in Figure 8.17. Except for two sharp peaks from the observed poaching data (blue line), the expert model (red line) seems to follow the same pattern as that of the observed poaching data. Comparing this with the outcomes of the tests does not shed much light on whether or not the data could in fact have resulted from the expert model or the uniform model. The only conclusion that can be reached is that the poaching data might have come from these two models, but that the evidence is inconclusive. Three possible explanations for this is (1) the large number of degrees of freedom, (2) the sparse data, and (3) the fact that the regions are too large thereby combining too many features in one region and “averaging out” over these features. A possible way to mitigate the problem with too-small cells and too-large regions, is to cluster cells together that share the same geographical features and use those clusters as regions.

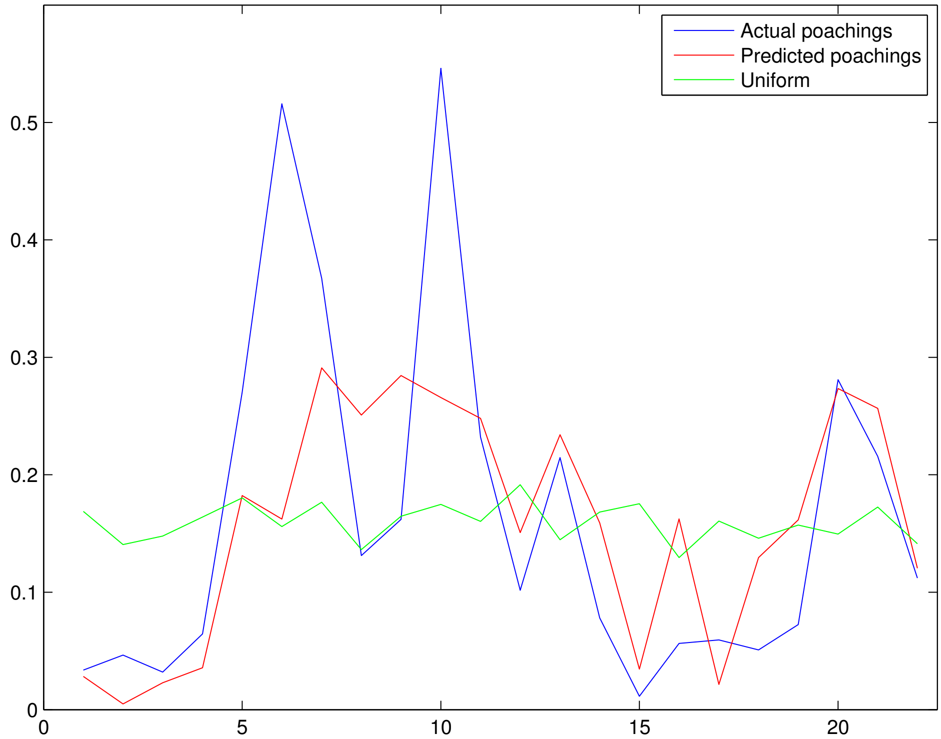

Figure 8.17. Goodness-of-fit test for observed, expected, and uniform values

8.6 CONCLUSION

This chapter presents a range of new insights that were garnered and changes that were made in order to align the model better with the experts' belief. A number of what-if analyses were performed to test whether or not the model agrees with logic and the knowledge of the experts. The model was shown to experts in the KNP and they agreed with most of the results. Minor changes were suggested to some of the states of the nodes, but mostly the outcomes of the what-if analyses correspond to the experts' belief. Some changes that were suggested concern the weighting given to certain states. This is an interesting result, as this reinforces our belief that the system is adaptive. Two years ago certain states were considered more important than others, and after this expert evaluation, it came to light that different states of the same variables are now considered more important. This also underlines the importance of updating the model frequently.

There are numerous ways of validating a model, some of which have been explained in this section. A more traditional way of model validation is to perform validation with data. This was performed by means of a goodness-of-fit test on both the predicted poaching values resulting from the expert model, as well as the uniform case where each cell has an equal probability of having a poaching event. The results showed that the observed poaching events could have come from both the expert model and the uniform model, but that it is more likely that it came from the expert model. Comparing this visually does not really instill much faith in the results of the goodness-of-fit tests, thus we conclude that, although the null hypotheses could not be rejected, there is insufficient evidence to indicate that they should be accepted.

CHAPTER 9 CONCLUSION

9.1 INTRODUCTION

The previous chapters detailed the rhino poaching problem and laid out all the tools that are needed to approach this complex problem. The rhino poaching model progressed from a first-order model to one that presents a realistic picture of the situation. This chapter concludes the thesis by summarising the most important findings and contributions. The reader is reminded of the research objectives presented at the start of the thesis and the results are discussed.

The research questions for this study, derived from the objectives in Chapter 1, were:

- Can the area be reduced that game rangers have to patrol?
- Is it possible to predict the area where poaching events are more probable?
- Can a decision support tool help increase the survivability of rhinos (where *survivability* refers to more pre-emptive strikes (rangers) or less poaching strikes (poachers))?
- Can a transdisciplinary, co-creation approach be facilitated?
- Is it possible to shift the ill-structured rhino poaching problem to a space where it is better understood?

9.2 SUMMARY OF FINDINGS

The first two research questions can be answered together. The output of the model is a heat map that highlights areas with a high probability of poaching attacks for specific time periods. The first research question can be answered: Yes, the area that game rangers have to patrol can be reduced. It would be too difficult to reduce the patrol area to the metre because there are many external factors such as the coarse time scale influencing the poachers as well as the rhinos. This adds uncertainty to the problem and makes prediction challenging. Thus to answer the second research question, it is possible to predict the general area where a poaching event will occur, but more data are needed to verify the accuracy of the predictions.

The latest version of the model was recently used by a group of experts to reason about the rhino poaching problem, and also the survivability of the rhinos. The experts could use this tool to test scenarios and discuss alternative management strategies. Although the model has not yet been tested in the field, its strategic capabilities have been explored. The model is currently used as a reasoning and decision support tool.

In the course of this study it became clear that collaboration is important. Through a transdisciplinary approach the researcher used the model as a vehicle for reasoning and for bringing together experts and stakeholders from many different fields and departments. This also leads to the last research question. At the start of this study not much was known about the rhino poaching problem and through this study we managed to shift the rhino poaching problem to a space where researchers and stakeholders can start to reason about the problem in a structured manner. It is now possible to see how the trends and the local drivers behind poaching shifted, and the experts concur that the insights from the model confirm theories that they have had for years but had no real evidence for.

9.3 SUMMARY OF CONTRIBUTIONS

The contributions made in this study are as follows:

- This is a first-of-its-kind Bayesian Network (BN) model of the rhino poaching problem developed through the capture of expert opinion, insights from multidisciplinary literature, and common sense.

- A “prior perspective” BN was developed before going into the expert workshop: experts had to evaluate and validate the network instead of starting with a “clean slate” network. This showed that time could be saved on the initial stages of the project, as well as eliminating the availability constraint partly.

- This body of work contributes to the structuring and understanding of the problem of rhino poaching, by delivering a tool that captures the key aspects of the problem space.

- Prior discussions as well as the expert workshops provided key features necessary to predict a region where a next poaching event could take place.

- The model (and tool) provides a manner in which to reduce the area that rangers need to patrol. The rangers can be sent to areas in the park that have the highest poaching probability, then the areas with the second highest poaching probability, and so forth.

- The validity of the model is confirmed by the what-if analyses in Section 8.3 that correctly described best-case and worst-case scenarios for the poaching problem, according to the experts.

- Decision support is provided to decision makers by helping them to think about the problem in a systematic way, as well as to capture the understanding, and parameterising a model with that understanding.

- A Bayesian Network, expert knowledge, Geographical Information Systems (GIS), and systems engineering concepts are combined in a transdisciplinary manner for a real-world ecological problem. Systems engineering approaches (such as the spiral approach) were used to gather requirements for the model and to iteratively improve the model.

- The model is used as the artefact to enable discourse, shared discovery, and understanding of the rhino poaching problem space.

- The ill-structured rhino poaching problem was shifted to a space in our sphere of influence where researchers can work with an abstraction of the problem.

- It was discovered that rhinos are not the victims or targets, but rather the commodities of the problem. Previous work on predictive policing could not be used directly owing to the fact that predictive policing techniques assume that the subject of the study (victim) can adapt their behaviour.

- The author showed that the system is reflective and adaptive and much effort was spent to understand the causal dynamics of the problem. The systems engineering spiral approach was applied numerous times to capture the dynamic property of the causal relationships.

- Collaboration between multiple departments and research areas was facilitated (transdisciplinary approach). Departments include various departments from the Council for Scientific and Industrial Research (CSIR), South African National Parks (SANParks), *etcetera*, and the disciplines involved includes mathematics, statistics, engineering, ecology, and criminology.

- Arriving at a solution to a real-world ecological problem was achieved through multiple stakeholder interactions and workshops, demonstrating that collaboration is vital in these complex situations which involve transdisciplinary expertise.

- This study presents one of the only known applications of BNs to wildlife crime.

9.4 FUTURE WORK

9.4.1 Excel-Matlab® elicitation template

Future work includes automating the breakdown of large unmanageable CPTs into smaller CPTs. A way also needs to be found to insert the probabilities of these smaller CPTs into the Matlab® code that reads in the *CPT* sheet. Lastly, the BNT network needs to be exported to Hugin®, but this has been a challenge up to now due to the age of the available software.

9.4.2 The model

The BN developed for this study is not complete nor is it perfect. There will always be new insights into the model, whether related to data or expert opinion. This is especially prevalent in the model structure as has been seen numerous times in this thesis.

A few changes to the existing network is recommended, such as more accurate estimates of impenetrable vegetation and open areas (dominant woody cover). These can be obtained with the aid of dedicated vegetation maps. This will result in the *Vegetation_density* node giving a more accurate view of the park in terms of the vegetation density.

Another change is to include secondary access roads and foot paths in the calculation of *Proximity_to_static_deterrents*, and to obtain real-time incursion point maps to update the *Distance_to_incursion_points* node. Data such as ranger patrol routes could be added to the *Ranger_present* node that will add a spatial factor to the *Active_deterrents* subgroup.

The home ranges of the white rhinos in the Kruger National Park (KNP) can be analysed to calculate what the optimal cell size should be for the grid. The same calculation could be performed using the KNP's road network.

Another modification to the model could be to the *Water* variable. It would be interesting to have the type of water source identified, for example if it is a permanent or seasonal water source. Input into this node could then be time of year, rainfall, and temperature. This might be useful in predicting a next poaching location.

A suggestion made by the experts is that a "return period" needs to be calculated as it is known that poachers return to the area where they poached rhinos, thereby wiping out an entire group in the course of a few weeks or months. Another suggestion was to speak to apprehended poachers and ask them why they do not poach certain rhinos, or do not poach at certain locations. Another feature that was deemed important was cellphone signals. This could be used to triangulate the location of a poacher, or at least, suspicious behaviour.

In this study one formal expert workshop was held after which multiple follow-up sessions were conducted with individual experts. Possible future work could see more formal expert workshops where each group of experts evaluates an iteration of the model until it clearly converges to a single structure.

A very important piece of future research is to alter the expert knowledge model so that it can be trained on data and thereby be used for prediction. Discussions have already been held as to how the model can be altered to keep the structure and the semantics more or less the same, but to change it so that it can use available data. These discussions were with colleagues who are working on situation awareness software with adaptive temporal capability specifically for the rhino poaching problem and who can supply the author with the necessary data, as well as testing opportunities. Since many poaching events have occurred during the course of this study, poaching data can and will be used to evaluate this model as well as other models that are based purely on poaching data, as opposed to expert knowledge.

9.4.3 Future applications

Future applications include integrating the prediction model into decision support tools for visualising and predictive capabilities. The model could also be generalised to include applications to other domains. This model is park-specific: for smaller parks there are other important factors, and the cell size will need to be adapted. The model could be adapted to work in any size park, and for the poaching of any type of animal. The goal is to have a tool that can be altered to suit any situation, be it lion poaching in KwaZulu-Natal, or abalone poaching in the Cape. Another application that is worth pursuing is that of neighbourhood crime. Many people are victims to crime every day and currently no such prediction model exists in South Africa. The idea is to develop a pro-active tool to predict areas of future neighbourhood crime.

9.5 CONCLUSION

Rhino poaching has reached epidemic status in South Africa where 90% of the world's rhino population resides. Most methods that are tried are reactive because when a rhino is poached, a task team sets off to capture the poacher. This study proposes a pre-emptive model to predict the area of a next poaching event so as to maximise the efficiency of the park's resources.

At the start of this study, little was known about the rhino poaching problem except what was seen on the news and experienced in the field. Literature only concentrated on the economic impact of the rhino horn trade or the ecological challenges facing the world if the rhino population were to become extinct. Even though rhinos were being poached by the hundreds, there were no complete high quality rhino poaching data sets to use. This study uses one of the most important sources of information, namely expert opinion to build and refine a model of the rhino poaching problem. Using data was always the goal, but in order to at least start working on the problem, expert knowledge was used. A prior model was developed by BN experts after which a group of domain experts were called together to evaluate the model. Changes were made, and after consulting with more experts on a one-on-one basis, consensus was reached on what the structure and the semantics of the model should be. Using BNs afforded us the option of using only expert knowledge, with the idea of later combining that expert knowledge with empirical data without necessitating the reworking of the entire model.

During the study it came to light that the current way of viewing rhino poaching might be incorrect. Rhinos are seen as the victims who are killed for their horns, when in actual fact they are the commodities with which syndicates and cartels launder money, smuggle drugs, buy weapons, and participate in human trafficking. Thinking of the problem this way highlights the fact that we might be collecting rhino poaching data incorrectly. Instead of only capturing the date and location, it might be of vital importance to also capture socio-economic trends outside the park. Did anything important happen in world affairs? Was a drug ring connected to smugglers in South Africa? What was the exchange rate of the South African rand to the Chinese yen, or the American dollar? Capturing these features might aid in detecting more patterns, and these patterns might lead to stopping poaching at the source, and not at the ground-level.

Another aspect that was highlighted is the reflective and adaptive nature of the problem. As soon as a policy change in the KNP occurs, the entire poaching process seems to shift. For a long time the poachers preferred to poach during the full moon, and when that fact was established by the people in charge of safeguarding the rhinos, the policy was changed so that all of the rangers were sent to the poaching hot spots during full moon. This policy change necessitated a shift in the operations of the poachers, as they could no longer poach during full moon with reckless abandon. This is both good news and bad news for any predictive model. On the positive side, continuing to shift the poachers' operations might end in them either giving up entirely, or getting frustrated and desperate, and thus making mistakes. However, this creates a challenge for the predictive model in terms of training and testing it on existing data. This was exactly the problem we faced during this study: it was not known when which policy changes took effect, thus it created challenges in training a certain version of the model on certain data sets.

A better approach to the problem might have been to have more formal expert workshops with different groups of experts. A suggested idea was that a "clean slate" approach is followed with one expert workshop, while the other workshop would then continue with the prior perspective model as in this thesis. The two models could then be compared to see if the prior perspective model biased the outcome of the workshop in any way. Another aspect that could have been addressed differently is that of data. The model could have been developed in such a way that it would be easier to train it on available data, as opposed to the number of abstract variables that it currently contains, making training a challenge.

An important contribution of this work is the BN model and the graphical structure that can be used to reason about the rhino poaching problem. Currently the expert-driven model is not used to make predictions, but it is rather used to encourage discussions and decision making. During the last expert meeting, the experts acknowledged that this model could indeed help them to make decisions about the problem as well as highlighting patterns and relationships between variables that they did not see before. The model also provides opportunity for the experts to see what type of data should be gathered in the future.

Owing to the prediction tool, we now have a good idea of the drivers of the rhino poaching problem. Knowledge is captured and shared, and an understanding about the problem now exists that did not exist a few years ago. New knowledge can be included in the model by switching in new nodes and switching out other nodes.The main contribution is that a reasoning tool now exists that can be used to facilitate discussions surrounding the safeguarding of the rhino population by reducing the area that rangers need to patrol.

REFERENCES

[1] A. Modise, "Minister Edna Molewa Highlights Progress in the War Against Poaching and Plans for 2015," jan 2015. [Online]. Available: https://www.environment.gov.za/mediarelease/molewa_waragainstpoaching2015

[2] R. Duffy, R. Emslie, and M. Knight, "Rhino poaching: How do we respond?" in *Evidence on Demand*, London, UK, oct 2013, pp. 1–35. [Online]. Available: http://www.evidenceondemand.info/rhino-poaching-how-do-we-respond

[3] E. Lunstrum, "Green Militarization: Anti-Poaching Efforts and the Spatial Contours of Kruger National Park," *Annals of the Association of American Geographers*, vol. 104, no. 4, pp. 816–832, jun 2014. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/00045608.2014.912545

[4] A. Modise, "Update on Rhino Poaching Statistics," 2013. [Online]. Available: https://www.environment.gov.za/?q=content/updaterhinopoachin_statistics

[5] E. Lunstrum, "Articulated Sovereignty: Extending Mozambican State Power Through the Great Limpopo Transfrontier Park," *Political Geography*, vol. 36, pp. 1–11, 2013. [Online]. Available: http://dx.doi.org/10.1016/j.polgeo.2013.04.003

[6] S. M. Ferreira, C. Greaver, G. A. Knight, M. H. Knight, I. P. J. Smit, and D. Pienaar, "Disruption of rhino demography by poachers may lead to population declines in Kruger National Park, South Africa," *PLoS ONE*, vol. 10, no. 6, pp. 1–18, 2015.

[7] T. L. Hieronymus, L. M. Witmer, and R. C. Ridgely, “Structure of White Rhinoceros (Ceratotherium simum) Horn Investigated by X-ray Computed Tomography and Histology With Implications for Growth and External Form,” *Journal of Morphology*, vol. 267, pp. 1172–1176, 2006.

[8] M. Montesh, “Rhino Poaching: A New Form of Organised Crime,” Tech. Rep., 2012.

[9] R. Ellis, “Rhino Horn: Facts and Myths,” in *Tiger Bone & Rhino Horn: The Destruction of Wildlife for Traditional Chinese Medicine*. Washington: Island Press, 2005, pp. 1–312.

[10] K. Nowell, “Assessment of Rhino Horn as a Traditional Medicine,” TRAFFIC, Tech. Rep., 2012. [Online]. Available: https://cites.org/eng/com/sc/62/E62-47-02-A.pdf

[11] K. Conrad, “Trade Bans: A Perfect Storm for Poaching?” *Tropical Conservation Science*, vol. 5, no. 3, pp. 245–254, 2012. [Online]. Available: www.tropicalconservationscience.org

[12] C. Eloff, “Rhino Poaching in South Africa - Is it a Losing Battle?” *PositionIT*, vol. Jan/Feb, pp. 57–61, 2012.

[13] S. F. Pires and W. D. Moreto, “Preventing Wildlife Crimes: Solutions That Can Overcome the ‘Tragedy of the Commons’,” *European Journal on Criminal Policy and Research*, vol. 17, no. 2, pp. 101–123, feb 2011. [Online]. Available: http://link.springer.com/10.1007/s10610-011-9141-3

[14] M. Nsehe, “South African Rhino Poacher Gets 77 Years In Jail,” jul 2014. [Online]. Available: http://www.forbes.com/sites/mfonobongnsehe/2014/07/25/rhino-poacher-gets-77-year-jail-term/

[15] D. Biggs, F. Courchamp, R. Martin, and H. P. Possingham, “Legal Trade of Africa‘s Rhino Horns,” *Science*, vol. 339, no. March, pp. 1038–1039, 2013.

[16] R. Duffy and F. A. V. St John, “Poverty, Poaching and Trafficking: What are the Links?” London, UK, Tech. Rep. June, 2013. [Online]. Available:

http://eprints.soas.ac.uk/17836/1/EoD_HD059_Jun2013_Poverty_Poaching.pdf

[17] G. L. Warchol, "The Transnational Illegal Wildlife Trade," *Criminal Justice Studies*, vol. 17, no. 1, pp. 57–73, mar 2004. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/08884310420001679334

[18] M. Eustace, "Rhino Poaching: What is the Solution?" pp. 1–5, jan 2012. [Online]. Available: http://www.bdlive.co.za/articles/2012/01/20/rhino-poaching-what-is-the-solution

[19] S. Mouton, "SANParks Takes Fight to Poachers," pp. 1–2, 2012. [Online]. Available: http://www.timeslive.co.za/thetimes/2012/12/13/sanparks-takes-fight-to-poachers

[20] J. R. Muntifering, W. L. Linklater, S. G. Clark, S. !Uri-≠ Khob, J. K. Kasaona, K. /Uiseb, P. Du Preez, K. Kasaona, P. Beytell, J. Ketji, B. Hambo, M. A. Brown, C. Thouless, S. Jacobs, and A. T. Knight, "Harnessing values to save the rhinoceros: insights from Namibia," *Oryx*, pp. 1–8, 2015.

[21] S. Ferreira, M. Hofmeyr, D. Pienaar, and D. Cooper, "Chemical Horn Infusions: A Poaching Deterrent or an Unnecessary Deception?" *Pachyderm*, vol. 55, pp. 54–61, 2014.

[22] D. P. Gonçalves, "Addressing Large National Problems in South Africa," in *INCOSE EMEA Sector Systems Engineering Conference*, Somerset West, 2014.

[23] H. Roodt and H. Koen, "Engineering Design of an Environmental Management System: A Transdisciplinary Response to the Rhino Poaching Problem," in *INCOSE EMEA Sector Systems Engineering Conference*, Somerset West, 2014.

[24] H. Mouton and J. P. De Villiers, "A Prologue to Estimating the Intent of a Potential Rhino Poacher," *Journal of Information Warfare*, vol. 11, no. 3, 2012.

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

[26] R. B. Martin, "A Legal Trade in Rhino Horn: Hobson's Choice," 2011. [Online]. Available: http://www.the-eis.com/data/literature/Hobsons Choice.pdf

[27] A. Modise, "Minister Edna Molewa Highlights Progress in the Fight Against Rhino Poaching," 2016. [Online]. Available: https://www.environment.gov.za/mediarelease/molewa_highlightsprogress_againstrhinopoaching

[28] K. Overland Park, "Definition and Types of Crime Analysis (White Paper 2014-02)," Tech. Rep., 2014.

[29] M. T. Plehn, "Control Warfare: Inside the OODA Loop," M.S. thesis, Dept. of Advanced Airpower Studies, Air Uni., Maxwell Air Force Base, Alabama, 2000.

[30] V. Kumar, A. Holzkaemper, B. Surridge, P. Rockett, M. Niranjan, and D. Lerner, "Bayesian Challenges in Integrated Catchment Modelling," in *4th Biennial Meeting of iEMSs*, 2008, pp. 362–371. [Online]. Available: http://eprints.lancs.ac.uk/51621/

[31] W. M. Dlamini, "A Bayesian Belief Network Analysis of Factors Influencing Wildfire Occurrence in Swaziland," *Environmental Modelling & Software*, vol. 25, no. 2, pp. 199–208, feb 2010.

[32] S. Johnson and K. Mengersen, "Integrated Bayesian Network Framework for Modeling Complex Ecological Issues," *Integrated Environmental Assessment and Management*, vol. 8, no. 3, pp. 480–490, jul 2012. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/21853523

[33] S. Johnson, L. Marker, K. Mengersen, C. H. Gordon, J. Melzheimer, A. Schmidt-Küntzel, M. Nghikembua, E. Fabiano, J. Henghali, and B. Wachter, "Modeling the Viability of the Free-Ranging Cheetah Population in Namibia: An Object-Oriented Bayesian Network Approach," *Ecosphere*, vol. 4, no. July, pp. 1–19, 2013. [Online]. Available: http://www.esajournals.org/doi/abs/10.1890/ES12-00357.1

[34] S. K. Frank, "Expert-based Bayesian Network Modeling for Environmental Management," Institute of Physical Geography, Goethe Univ. Frankfurt, Frankfurt am Main, Germany, Tech.

Rep., 2015.

[35] L. R. Hope, A. E. Nicholson, and K. B. Korb, "Knowledge Engineering Tools for Probability Elicitation," Tech. Rep., 2002.

[36] S. Low Choy, R. O'Leary, and K. Mengersen, "Elicitation by Design in Ecology: Using Expert Opinion to Inform Priors for Bayesian Statistical Models," *Ecology*, vol. 90, no. 1, pp. 265–277, jan 2009. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/19294931

[37] M. Düspohl, S. Frank, and P. Döll, "A Review of Bayesian Networks as a Participatory Modeling Approach in Support of Sustainable Environmental Management," *Journal of Sustainable Development*, vol. 5, no. 12, pp. 1–18, nov 2012. [Online]. Available: doi: 10.5539/jsd.v5n12p1

[38] L. Uusitalo, "Advantages and Challenges of Bayesian Networks in Environmental Modelling," *Ecological Modelling*, vol. 203, no. 3-4, pp. 312–318, may 2007. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0304380006006089

[39] R. K. Mccann, B. G. Marcot, and R. Ellis, "Bayesian Belief Networks: Applications in Ecology and Natural Resource," *Canadian Journal of Forest Research*, vol. 36, pp. 3053–3062, 2006.

[40] A. M. Ellison, "An Introduction to Bayesian Inference for Ecological Research and Environmental Decision-Making," *Ecological Applications*, vol. 6, no. 4, pp. 1036–1046, 1996.

[41] A. De Waal and T. Ritchey, "Combining Morphological Analysis and Bayesian Networks for Strategic Decision Support," *ORiON*, vol. 23, no. 2, pp. 105–121, 2007.

[42] D. Barber, *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2015. [Online]. Available: http://web4.cs.ucl.ac.uk/staff/D.Barber/textbook/181115.pdf

[43] R. E. Neapolitan, *Learning Bayesian Networks*. New Jersey, USA: Prentice-Hall, Inc., 2003.

[44] R. Fisher, R. O'Leary, S. Low-Choy, K. Mengersen, and M. J. Caley, "A Software Tool for Elicitation of Expert Knowledge about Species Richness or Similar Counts," *Environmental Modelling and Software*, vol. 30, pp. 1–14, 2012.

[45] A. James, S. Low-Choy, and K. Mengersen, "Elicitator: An Expert Elicitation Tool for Regression in Ecology," *Environmental Modelling and Software*, vol. 25, no. 1, pp. 129–145, 2010. [Online]. Available: http://dx.doi.org/10.1016/j.envsoft.2009.07.003

[46] G. Van Houtven, C. Mansfield, D. J. Phaneuf, R. von Haefen, B. Milstead, M. A. Kenney, and K. H. Reckhow, "Combining Expert Elicitation and Stated Preference Methods to Value Ecosystem Services from Improved Lake Water Quality," *Ecological Economics*, vol. 99, pp. 40–52, mar 2014. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0921800914000032

[47] R. O'Leary, S. Low-Choy, J. V. Murray, M. Kynn, R. Denham, T. G. Martin, and K. Mengersen, "Comparison of Three Expert Elicitation Methods for Logistic Regression on Predicting the Presence of the Threatened Brush-Tailed Rock-Wallaby Petrogale Penicillata," *Environmetrics*, vol. 20, pp. 379–398, 2009. [Online]. Available: http://oro.open.ac.uk/7797/

[48] R. Denham and K. Mengersen, "Geographically Assisted Elicitation of Expert Opinion for Regression Models," *Bayesian Analysis*, vol. 2, no. 1, pp. 99–136, mar 2007. [Online]. Available: http://projecteuclid.org/euclid.ba/1340390065

[49] B. Das, "Generating Conditional Probabilities for Bayesian Networks: Easing the Knowledge Acquisition Problem," *CoRR*, pp. 1–24, 2004. [Online]. Available: http://arxiv.org/abs/cs/0411034

[50] H. Wang and M. J. Druzdzel, "User Interface Tools for Navigation in Conditional Probability Tables and Elicitation of Probabilities in Bayesian Networks," *Uncertainty in Artificial Intelligence Proceedings*, pp. 617–625, 2000. [Online]. Available: http://uai.sis.pitt.edu/papers/00/p617-wang.pdf

[51] J. Pellikka, S. Kuikka, H. Lindén, and O. Varis, "The role of game management in wildlife populations: uncertainty analysis of expert knowledge," *European Journal*

of Wildlife Research, vol. 51, no. 1, pp. 48–59, feb 2005. [Online]. Available: http://link.springer.com/10.1007/s10344-004-0073-9

[52] I. Albert, S. Donnet, C. Guihenneuc-Jouyaux, S. Low-Choy, K. Mengersen, and J. Rousseau, "Combining Expert Opinions in Prior Elicitation," *Bayesian Analysis*, vol. 7, no. 3, pp. 503–532, 2012.

[53] F. Hansson and S. Sjökvist, "Modelling Expert Judgement Into a Bayesian Belief Network: A Method for Consistent and Robust Determination of Conditional Probability Tables," Master's thesis, Centre for Mathematical Sciences and Mathematical Statistics, Lund Univ., Lund, Sweden, 2013.

[54] T. Kokkonen, H. Koivusalo, H. Laine, A. Jolma, and O. Varis, "A Method for Defining Conditional Probability Tables with Link Strength Parameters for a Bayesian Network," *International Congress on Modelling and Simulation*, pp. 428–434, 2005. [Online]. Available: http://www.mssanz.org.au/modsim05/papers/kokkonen.pdf

[55] C. Spaccasassi and L. Deleris, "A Web-Based Tool for Expert Elicitation in Distributed Teams," *CEUR Workshop Proceedings*, vol. 818, pp. 125–133, 2011.

[56] R. Richards, O. Sahin, M. Sano, J.-O. Meynecke, and R. Tiller, "App2Adapt: Using Tablet Technology to Elicit Conditional Probabilities for Bayesian Belief Network Modelling," in *7th Inl. Congress on Environmental Modelling and Software*, 2014.

[57] D. V. Pullar and T. H. Phan, "Using a Bayesian Network in a GIS to Model Relationships and Threats to Koala Populations Close to Urban Environments," *Modsim 2007: International Congress on Modelling and Simulation: Land, Water and Environmental Management: Integrated Systems for Sustainability*, pp. 1370–1375, 2007.

[58] V. Stelzenmüller, J. Lee, E. Garnacho, and S. I. Rogers, "Assessment of a Bayesian Belief Network-GIS Framework as a Practical Tool to Support Marine Planning," *Marine pollution bulletin*, vol. 60, no. 10, pp. 1743–54, oct 2010. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/20619861

[59] S. Johnson, K. Mengersen, A. de Waal, K. Marnewick, D. Cilliers, A. M. Houser, and L. Boast, "Modelling Cheetah Relocation Success in Southern Africa Using an Iterative Bayesian Network Development Cycle," *Ecological Modelling*, vol. 221, no. 4, pp. 641–651, feb 2010. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0304380009007947

[60] R. Farmani, H. J. Henriksen, D. Savic, and D. Butler, "An Evolutionary Bayesian Belief Network Methodology for Participatory Decision Making Under Uncertainty: An Application to Groundwater Management," *Integrated Environmental Assessment and Management*, vol. 8, no. 3, pp. 456–461, jul 2012. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/21412969

[61] S. H. Hamilton, C. A. Pollino, and A. J. Jakeman, "Habitat Suitability Modelling of Rare Species Using Bayesian Networks: Model Evaluation Under Limited Data," *Ecological Modelling*, vol. 299, pp. 64–78, 2015. [Online]. Available: http://dx.doi.org/10.1016/j.ecolmodel.2014.12.004 http://linkinghub.elsevier.com/retrieve/pii/S0304380014006103

[62] K. Baumgartner, S. Ferrari, and G. Palermo, "Constructing Bayesian Networks for Criminal Profiling from Limited Data," *Knowledge-Based Systems*, vol. 21, no. 7, pp. 563–572, 2008.

[63] C. Zhang, A. Sinha, and M. Tambe, "Keeping Pace with Criminals: Designing Patrol Allocation Against Adaptive Opportunistic Criminals," in *Proc. 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, E. Bordini and Y. Weiss, Eds., no. 1, Istanbul, Turkey, 2015, pp. 1351–1359.

[64] G. Hoskin, "Major Boost in Resources to Fight Poachers," 2011. [Online]. Available: http://www.iol.co.za/news/crime-courts/major-boost-in-resources-to-fight-poachers-1103126

[65] E. J. Milner-Gulland and N. Leader-Williams, "A Model of Incentives for the Illegal Exploitation of Black Rhinos and Elephants: Poaching Pays in Luangwa Valley, Zambia," *The Journal of Applied Ecology*, vol. 29, no. 2, pp. 388 – 401, 1992. [Online]. Available: http://www.jstor.org/stable/2404508?origin=crossref

[66] K. Leggett, "Coprophagy and Unusual Thermoregulatory Behaviour in Desert-Dwelling

Elephants of North-Western Namibia," *Pachyderm*, vol. 36, pp. 113–115, 2004. [Online]. Available: http://iucn.org/afesg/pachy/pachy36.html

[67] E. Martin, "Rhino Poaching in Nepal During an Insurgency," *Pachyderm*, vol. 36, pp. 87–98, 2004. [Online]. Available: http://iucn.org/afesg/pachy/pachy36.html

[68] C. Lockwood, "Analysis of Rhino Poaching in KwaZulu-Natal, South Africa," M.S thesis, Imperial College London, 2010.

[69] J. S. Kahler and M. L. Gore, "Beyond the Cooking Pot and Pocket Book: Factors Influencing Noncompliance With Wildlife Poaching Rules," *International Journal of Comparative and Applied Criminal Justice*, vol. 36, no. 2, pp. 103–120, may 2012. [Online]. Available: http://www.tandfonline.com/doi/abs/10.1080/01924036.2012.669913

[70] J. M. Rowcliffe, E. de Merode, and G. Cowlishaw, "Do Wildlife Laws Work? Species Protection and the Application of a Prey Choice Model to Poaching Decisions." *Proceedings. Biological sciences / The Royal Society*, vol. 271, no. 1557, pp. 2631–2636, 2004.

[71] N. Park, E. Serra, T. Snitch, and V. S. Subrahmanian, "APE: A Data-Driven, Behavioral Model Based Anti-Poaching Engine," *IEEE Transactions on Computational Social Systems*, vol. 2, no. 2, pp. 15–37, 2015.

[72] R. Critchlow, A. J. Plumptre, M. Driciru, A. Rwetsiba, E. J. Stokes, C. Tumwesigye, F. Wanyama, and C. M. Beale, "Spatiotemporal Trends of Illegal Activities from Ranger-Collected Data in a Ugandan National Park," *Conservation Biology*, vol. 29, no. 5, pp. 1458–1470, 2015.

[73] C. M. Bishop, "Pattern Recognition and Machine Learning," in *Information Science and Statistics*, 1st ed., M. Jordan, J. Kleinberg, and B. Scholkopf, Eds. Springer-Verlag New York, 2006, pp. 1–738.

[74] S. Ross, *A First Course in Probability*, 6th ed. Prentice Hall, 2001.

[75] K. B. Korb and A. E. Nicholson, *Bayesian Artificial Intelligence, Second Edition*, 2nd ed. Florida, USA: CRC Press, Inc., 2010.

[76] S. L. Morgan and C. Winship, *Counterfactuals and Causal Inference: Methods and Principles for Social Research*, 2nd ed. Cambridge University Press, 2014.

[77] J. Pearl, "The Logic of Structure-Based Counterfactuals," in *Causality*, 2nd ed. Cambridge University Press, 2009, ch. 7, pp. 201–257.

[78] M. Kynn, "The 'Heuristics and Biases' Bias in Expert Elicitation," *Journal of the Royal Statistical Society. Series A: Statistics in Society*, vol. 171, pp. 239–264, 2008.

[79] M. E. Borsuk, S. Schweizer, and P. Reichert, "A Bayesian Network Model for Integrative River Rehabilitation Planning and Management," *Integrated Environmental Assessment and Management*, vol. 8, no. 3, pp. 462–472, jul 2012. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/21608112

[80] H. Wang, "Building Bayesian Networks: Elicitation, Evaluation, and Learning," Ph.D. dissertation, Faculty of Arts and Sciences, Univ. of Pittsburgh, Pittsburgh, 2004.

[81] L. E. Cohen and M. Felson, "Social Change and Crime Rate Trends: A Routine Activity Approach," *American Sociological Review*, vol. 44, no. 4, pp. 588–608, 1979.

[82] B. W. Boehm, "A Spiral Model of Software Development and Enhancement," *IEEE Computer*, vol. 21, no. 5, pp. 61–72, 1988.

[83] J. J. Ryan, T. A. Mazzuchi, D. J. Ryan, J. Lopez de la Cruz, and R. Cooke, "Quantifying Information Security Risks Using Expert Judgment Elicitation," *Computers & Operations Research*, vol. 39, no. 4, pp. 774–784, apr 2012. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0305054810002893

[84] M. C. Runge, S. J. Converse, and J. E. Lyons, "Which Uncertainty? Using Expert Elicitation and Expected Value of Information to Design an Adaptive Program,"

Biological Conservation, vol. 144, no. 4, pp. 1214–1223, apr 2011. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S000632071000529X

[85] L. Finlay and B. Gough, "The Reflexive Journey: Mapping Multiple Routes," in *Reflexivity: a practical guide for researchers in health and social sciences*. Blackwell Science Ltd., 2008, pp. 3–20.

[86] E. Baker, V. Bosetti, K. E. Jenni, and E. C. Ricci, "Facing the Experts: Survey Mode and Expert Elicitation," 2014.

[87] S. Esbjörn-Hargens and M. Zimmerman, *Integral Ecology: Uniting Multiple Perspectives on the Natural World*. Integral Books, 2009.

[88] G. Bucini, "Woody Cover in African Savannas: Mapping Strategies and Ecological Insights at Regional and Continental Scales," Ph.D. dissertation, Ecology, Colorado State Univ., Fort Collins, Colorado, 2010.

[89] W. Trollope, "Ecological Effects of Fire in South African Savannas," in *Ecology of Tropical Savannas*, B. Huntley and B. Walker, Eds. Springer Berlin Heidelberg, 1982, pp. 292–306.

[90] B. Van Wyk and P. Van Wyk, *Field Guide to Trees of Southern Africa*. Struik Publishers (Pty) Limited, 1997.

[91] K. Murphy, "How to Use the Bayes Net Toolbox," pp. 1–28, 2007. [Online]. Available: http://bayesnet.github.io/bnt/docs/usage.html

[92] K. Shin and J. K. Hammond, "Stochastic Processes; Correlation Functions and Spectra," in *Fundamentals of Signal Processing for Sounds and Vibration Engineers*, 2008, pp. 232–233.

[93] W. P. D. Gertenbach, "Landscapes of the Kruger National Park," *Koedoe - African Protected Area Conservation and Science*, vol. 26, no. 1, pp. 9–121, 1983. [Online]. Available: http://www.koedoe.co.za/index.php/koedoe/article/viewFile/591/620

[94] "Kruger National Park: General Information," 2016. [Online]. Available: https://www.sanparks.org/parks/kruger/tourism/general.php

[95] D. Pienaar, "Habitat Preference of the White Rhino in the Kruger National Park," in *Symposium on Rhinos as Game Ranch Animals*, 1994, pp. 59–64.

[96] D. Pienaar, J. d. P. Bothma, and G. Theron, "Landscape Preferences of the White Rhinoceros in the Southern Kruger National Park," *Koedoe*, vol. 35, no. 1, pp. 1–7, 1992.

[97] D. Pienaar, J. d. P. Bothma, and G. K. Theron, "Landscape Preference of the White Rhinoceros in the Central and Northern Kruger National Park," *Koedoe*, vol. 36, no. 1, pp. 79–85, 1993.

[98] "MATLAB - The Language of Technical Computing." [Online]. Available: http://www.mathworks.com/products/matlab/

[99] "Hugin Expert - The Leading Decision Support Tool." [Online]. Available: http://www.hugin.com/

[100] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 1st ed. Prentice-Hall, Inc., 1995.

[101] A. Kahn, "Topological Sorting of Large Networks," *Communications of the ACM*, vol. 5, no. 11, pp. 558–562, 1962.

[102] D. J. Pearce and P. H. J. Kelly, "A Dynamic Topological Sort Algorithm for Directed Acyclic Graphs," *Journal of Experimental Algorithmics*, vol. 11, no. 1, pp. 1–24, 2007.

[103] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, "Section 23.4: Topological Sort," in *Introduction To Algorithms*, 2nd ed. MIT Press and McGraw-Hill, 2001, ch. 22.4, pp. 485–488.

[104] G. Wang, S. Chen, X. Zhang, and Z. Zhou, "Topological Ordering Algorithm for LDAG," *Information Processing Letters*, vol. 112, no. 21, pp. 823–828, 2012. [Online]. Available: http://dx.doi.org/10.1016/j.ipl.2012.07.015

[105] C. Pang, J. Wang, Y. Cheng, H. Zhang, and T. Li, "Topological Sorts on DAGs," *Information Processing Letters*, vol. 115, no. 2, pp. 298–301, 2015. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0020019014002142

[106] H. Koen, J. De Villiers, H. Roodt, and A. De Waal, "An expert-driven causal model of the rhino poaching problem," *Ecological Modelling*, vol. 347, pp. 29–39, 2017. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0304380016307621

[107] K. Carley, "Validating computational models," 1996. [Online]. Available: http://www.econ.iastate.edu/tesfatsi/EmpValid.Carley.pdf

[108] R. Sargent, "Verification and validation of simulation models," in *Proceedings of the 1998 Winter Simulation Conference*, D. Medeiros, E. Watson, J. Carson, and M. Manivannan, Eds., 2005, pp. 121–130.

[109] D. Wackerly, W. Mendenhall, and R. Scheaffer, *Mathematical Statistics with Applications*, 6th ed. Duxbury Press, 2001.