Modelling uncertain adaptive decisions: Application to KwaZulu-Natal sugarcane growers

C.S. Price[1,4[0000-0002-8749-7253]], D. Moodley[2,4[0000-0002-4340-9178]] and A.W. Pillay[3,4[0000-0001-7160-6972]]

[1] University of KwaZulu-Natal, Westville, 4001, RSA
pricec@ukzn.ac.za
[2] University of Cape Town, Rondebosch, 7701, RSA
deshen@cs.uct.ac.za
[3] University of KwaZulu-Natal, Westville, 4001, RSA
pillayw4@ukzn.ac.za
[4] CSIR Meraka Centre for Artificial Intelligence Research, Pretoria, 0001, RSA

Abstract. A dynamic Bayesian decision network was developed to model the pre-harvest burning decision-making processes of sugarcane growers in a KwaZulu-Natal sugarcane supply chain and extends previous work by Price *et al.* (2018). This model was created using an iterative development approach. This paper recounts the development and validation process of the third version of the model. The model was validated using Pitchforth and Mengersen (2013)'s framework for validating expert elicited Bayesian networks. During this process, growers and cane supply members assessed the model in a focus group by executing the model, and reviewing the results of a pre-run scenario. The participants were generally positive about how the model represented their decision-making processes. However, they identified some issues that could be addressed in the next iteration. Dynamic Bayesian decision networks offer a promising approach to modelling adaptive decisions in uncertain conditions. This model can be used to simulate the cognitive mechanism for a grower agent in a simulation of a sugarcane supply chain.

Keywords: Dynamic Bayesian decision network, model development, expert model validation

1 Introduction

South African sugarcane growers face many uncertainties when delivering their cane to the mill. The main uncertainty is the weather. Before harvesting their cane to deliver it to the mill, growers generally burn it [1] to reduce transport costs. Rain interrupts burning and sometimes harvesting activities [2]. Long periods of soft soaking drizzle leave the cane wet, unable to be burned [3]. High winds pose a risk of runaway fires during burning [3]. If they believe that a period of adverse weather is forecast, growers could harvest more cane [2]. However, they cannot burn too much cane in advance: from the time it is burned, the cane quality starts to decline. As they are paid according to the quality of their cane, growers are loathe to deliver older cane to the mill. In addition, each grower has to deliver a consistent amount of cane to the mill during the

approximately 38 week milling season, and can face penalties if this is not done. The mill faces mechanical problems if it is starved of cane. The decision to burn cane is therefore a complex one. The grower faces the uncertainties of the weather; what he believes the weather will be, given the forecast; and his belief of the cane's dryness. He needs to be adaptable to burn when conditions are suitable and to burn in advance if needed. He also has the requirement of making regular daily deliveries and not falling behind. The aim of this study is to develop a model of the grower's pre-harvest burning decision making.

Bayesian networks (BNs) offer a compact, graphical way of modelling such problems. BNs explicitly represent uncertainty [4] and can be used to model complex domains [4]. They can be created from data, expert opinion, or a combination of both [5].

A BN was developed to model the sugarcane grower's pre-harvest burning decisions for a KwaZulu-Natal sugar mill's supply chain [6] (version 1); this paper extends that model. Version 1 took account of the grower's belief of the weather and cane dryness. It also made the decision of how much cane to burn adaptively, depending on how much cane was already burned, and what the grower's delivery status was (behind/on target/ahead). An assumption of this model was that if one burned cane, it could be delivered to the mill on the same day. Another was that if the grower were behind with deliveries, he could burn damp cane. If rain was forecast later that day, the grower could still burn (and deliver). Burning cane in high winds was not allowed. In version 2, advance weather forecasts were incorporated. Three summary nodes were added for the grower's belief of burning conditions, the advance and current day's burning requirements. Some assumptions were also changed: if one burned on one day, one could only deliver the next day; if rain was forecast that day, one could not burn. Damp cane had a low probability of being burned. Large scale growers burn and deliver daily, while small scale growers burn every two days and deliver daily. Version 2 was shown to the Mill's Cane Procurement Managers, who gave feedback and inputs. They highlighted that the mill can slow or accelerate deliveries by calling for less or more cane. The grower may also have a limited delivery capacity to catch up if he fell behind. These features were incorporated into version 3. Version 3 needed to be tested and validated by those experienced in cane growing and delivery. That process, and the model's development is reported on in this paper.

The layout of the paper is as follows: the literature review introduces BNs, their development and validation. The methodology section outlines how the model was developed and validated. Version 3 of the BN model is presented next, followed by its validation. A conclusion and proposals for future work follows. The focus group methodology, model's assumptions and larger figures can be found in the Appendices.

2 Literature review

2.1 Bayesian networks

Bayesian networks (BNs) are acyclic graphs which represent cause-effect relationships under uncertainty [5]. Variables are represented by nodes, which are linked to other variables by arcs. Nodes have a number of states representing the discretized range of

values that the node can have. A node also has a conditional probability table (CPT) which contains the probability of being in each of the nodes' states, given the parent node's states. If there is no parent node, *a priori* probability values are entered. When the BN is run, the user can choose a particular state of a node (called "entering evidence") [5]; this will set the value of that node's state to 100%. This action will update the probabilities of the states in the other nodes in the network (called belief updating).

There are four main types of Bayesian networks. Those which only consist of chance nodes are called Bayesian networks (BN) or Bayesian belief networks. Including a decision node (and utility node) turns the BN into a Bayesian decision network (BDN). These two types of network cannot represent how a system changes over time. A BN which is repeated in many time slices, with arrows linking nodes in the different slices is called a dynamic Bayesian network (DBN) [5]. A DBN which includes a decision in each time slice is called a dynamic Bayesian decision network (DBDN). However, not many DBDNs are reported in the literature, as they are not explicitly supported in most BN software packages. In the sugar domain, Drury *et al.* [7] generated a BN of factors affecting sugarcane yield using text mining. Price *et al.* [6] developed a DBDN to model how much cane growers decided to burn before harvesting. To our knowledge, no other models have been developed to represent these types of adaptive, uncertain decisions in the sugarcane supply chain.

2.2 BN model development

In their paper giving guidelines for developing BNs, Marcot *et al.* [8] propose using an iterative approach. They advise BN developers to keep the number of parent nodes to three or fewer so that the CPT size can be kept small enough to understand. They recommend that there should not be more than four nodes between the input node and the output node. They also advise developers to keep the number of states per node small enough to be able to produce the results needed by the model.

2.3 BN model validation

Ideally, one would prefer to validate a BN with data [5, 9]. The data used to generate the structure and/or CPTs needs to be split into two parts, i.e. one part to create the model and the other part to test it [5]. However, if data are not available, one must rely on experts to validate the network [9]. It is not sufficient to test only model outputs in one's validation. Pitchforth and Mengersen [9] proposed a framework of how to validate a Bayesian network based on experts' opinion. This framework has been used by several authors to validate their BN using experts (e.g. [10-12]). When using experts, care should be taken not to use the same experts for developing the Bayesian network model and for validating the model [9]. The framework gives the following categories of validity to consider when validating an expert elicited BN [9]:

- **Nomological validity** is an overarching category which considers how the model fits within the model domain's literature. It says little about the model's structure and behavior [10]. Pitchforth *et al.* [10] plot this type of validity by comparing the

model to be validated with other models in the literature and rating four aspects: model structure, discretization, parameterization and behavior.

- **Face validity** reflects whether the model's structure looks like the experts and literature expect; whether the states have been discretized in ways that the experts expect, and that the parameters (CPT values) are what the experts would expect.
- **Content validity** means that all the nodes and states that are needed in the model to produce the output are present, and no unnecessary nodes and states are present.
- **Concurrent validity** looks at whether the network or a section of the network behaves in an identical way to a section in another network.
- **Convergent validity** investigates how similar the structure of the model is to models from a similar domain; it also compares how the states are discretized and the CPT values with other models from a similar domain. It can also examine whether similar concepts in the same model are modelled in the same kind of way [10]. This also includes how the nodes states are discretized.
- **Discriminant validity** means that the model describing one domain should be different (in structure, discretization and parameterization) from a model of a different domain.
- **Predictive validity** considers both model behavior and model output. There are three main ways to test predictive validity:
  - *Behavior sensitivity tests.* Sensitivity to evidence examines how sensitive the network is to small changes in input; sensitivity to parameters examines how sensitive the network is to slight changes in the parameter (CPT) values. For the latter, one can start with a "sensitivity set" [13] of all nodes whose parameter changes would affect a particular node [5]. In BDNs, a Value of Information analysis is used to show what additional information would improve the decision [5].
  - *Extreme conditions tests.* These tests look at how node pairs, sub-networks or the network behaves when inputting parameters for extreme conditions.
  - *Qualitative features tests.* These tests evaluate the model's output. If other models are available, the output of the BN can be compared to that of other models. Extreme conditions tests are a special type of qualitative features test.

3 Methodology

3.1 BN development

The mill's Cane Procurement Manager and two large scale growers were interviewed[1] to discover business rules for how deliveries to the mill occur, and how cane burning and harvesting affect deliveries. The aim of the interviews was to develop an agent-based model of the sugarcane supply chain [14]. The content of the interviews was also used as the basis for developing the DBDN. The rationale for choosing a DBDN type of Bayesian network model is covered in [6].

BN software was reviewed to see if it could support DBDNs. However, no software was found with the capability of modelling a DBDN. In Hugin [15] ver 7.8, the Object-

[1] Ethical clearance HSS/0204/101

oriented modelling feature used for implementing DBNs did not have the capability to support the inter slice arrows that were needed in the model. It was therefore decided to build a model by copying the time slices and linking the relevant nodes manually with inter slice arrows. The model was developed iteratively using the knowledge engineering approach to developing BNs [5]. Each part of the model was tested by the author before adding more functionality to the model [16].

Actual weather data (1 January 1998 to 31 December 2015) from the mill area was used to populate the CPTs of the "actual weather" nodes. Based on interviews with the Cane Procurement Manager, actual rain for April-November was classified as drizzle if the daily total was ≤4mm and thundershowers if the total was >4mm; for December-March, half the daily rainfall (0.1-4mm) was assumed to be thundershowers. The data was imported into Netica [17] using the Expectation Maximisation algorithm. The resultant network was imported into Hugin [15] ver 7.8. This was later upgraded to ver 8.6. For more detail on this section of the network, see [6].

The structure of the rest of the DBDN was worked out manually. CPT values were entered into an Excel spreadsheet by the first author, based on business rules, using an anchoring technique [5] to compare different options, before being copied into Hugin.

3.2 BN model validation

Pitchforth and Mengersen [9]'s validation framework was used to validate the network. During the model's development, as each node's CPT was added, it was tested before moving on to the next node [16]. As the number of nodes with populated CPTs increased, the number of nodes tested increased. Model results were tested on a two time slice version of the model before adding further time slices.

Sensitivity tests and Value of Information analyses were run in Hugin on Day 1 nodes to test the sensitivity of the nodes to parameters in another node, and to determine which additional information would add to the Day 1 decision [5]. To test model outcomes, the model was tested programmatically by setting evidence in various nodes using the Hugin Application Programming Interface (API) and writing the model node states and results to a text file which was later imported into MS Excel. These runs were analysed to test extreme model conditions and model results.

Focus group meeting

A focus group meeting was organized to get feedback about the model. The Cane Procurement Manager invited five participants, all of whom had growing experience (see Table 1). Apart from the Cane Procurement Manger, none had heard of the model before. On the day of the meeting, the weather was cold, with drizzle.

After all had signed consent letters to participate in the focus group[2], the moderator (first author) gave an overview of the DBDN: its nodes, states and assumptions. At this point, participant 5 left. The DBDN was then demonstrated in Hugin. The participants asked the moderator to enter that and the following day's weather into the model to see how it would react to cold, drizzly weather and wet cane.

[2] Ethical clearance HSS/0204/101

Table 1. Focus group participants and their expertise

Participant No.	Growing Experience	Cane Supply Experience
1	Grew up on sugarcane farm; large scale grower (10 years)	
2	Owned farm (5 years)	5 years
3	Farm manager (2 years)	3 years
4	Grew up on sugarcane farm; farm manager (1 year)	5 years
5	Farm manager (5 years)	15 years

The results of a scenario which had previously been run were presented. The scenario was based on actual mill weather data from 5-13 October 2016. October 2016 was particularly wet, and since there had been much rain at that time, it was assumed that the cane was already damp and the grower was behind with deliveries. This was the same scenario as Scenario 5 from Price *et al.* [6], and was chosen because spring rain causes problems with deliveries. At the end of the meeting, participants were asked to participate in a "write down" [18] to gauge each participant's view independently of others' views. They briefly answered questions on the model's assumptions, and what aspects of the model needed changing. They were asked to rate the model's representativeness of what happens in the mill's supply chain on a five point Likert scale question (Very poorly to Very well). Finally, they were asked for additional comments on the model. More details of the focus group methodology can be found in Appendix A.

4 The dynamic Bayesian decision network

Version 3 of the DBDN developed to represent growers' pre-harvest burning decisions is shown in Fig. 1[3]. Two days (Day 0 and Day 1) are shown in this image. Day 1's time slice is replicated for further days (see Appendix B for the model assumptions.) There are four categories of nodes in this DBDN: "Today's weather" (top left of Day 1) models today's maximum temperature, rain and wind forecasts, and the grower's belief of the forecasts. The "Advance burning requirements" nodes (top right of Day 1) models the grower's belief in high winds or rain within the next seven days. The "Dryness of cane" nodes spans two time slices. Previous day's actual weather affects the grower's belief in the cane's dryness on that day. Finally, the "Delivery" nodes (bottom of the model) also span two time slices. The cane already burned on any day is affected by how much cane was burned the previous day and how much was delivered. The status of deliveries depends on the previous day's status, and how much was delivered the previous day. Because the mill's call for cane is modelled in increments of 0.25 day's quota, the model decides how much cane to burn and deliver in 0.25 day increments.

[3] See Fig. 3 in Appendix B for a larger version.

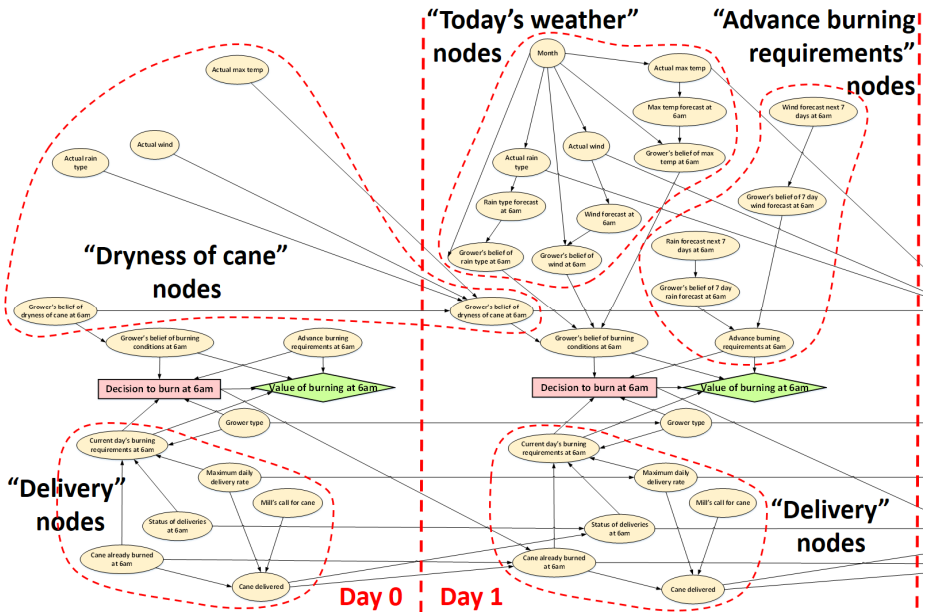

Fig. 1. The dynamic Bayesian decision network which represents the grower's decision to burn cane before harvesting. Day 1's nodes are replicated for subsequent days.[4]

5 Model validation results

The validation of the DBDN, version 3, using Pitchforth and Mengersen [9]'s framework, is outlined below.

5.1 Nomological validity

BNs relating to the sugar domain are scarce. Drury *et al.* [7] generated a BN of factors affecting sugarcane yield using text mining. In a farming context, Clemen [19] showed the structure of a farmer's sequential decision to protect his trees from adverse weather. This model showed one decision per day in a time slice, which is similar in concept to the DBDN presented here. Price *et al.* [6] developed a DBDN of grower's pre-harvest cane burning decisions. This work also modelled how much cane could be delivered to the mill. Stray *et al.* [20] used an optimization model to decide which cane field to harvest at a tactical level. A simulation model was found which model the delivery of cane to the mill on a weekly basis [21], on an hourly basis [22] and on a minute basis [14]. However, these models did not consider the burning of cane. Apart from [6], no models were found which represent sugarcane growers' pre-harvest burning decisions.

[4] See Fig. 3 in Appendix B for a larger version.

5.2 Face validity

During model development, the second author reviewed the model, and made suggestions for improvements. During the meetings with the supply chain members to review versions 2 and 3, the participants did not question the presence of any nodes. In the focus group meeting, the participants said that they would not burn a fraction of a day's quota; they would rather delay and burn slightly more, so the three states ('burn 0.25 days' cane'; 'burn 0.5 days' cane'; 'burn 0.75 days' cane') could be removed from the model. They also said that large growers burn every other day (three times a week). In the model, this is how the "small growers" work. In reality, small growers only burn about twice a week. (This is contrary to what was explained by the Cane Procurement Manager in the version 2 review).

5.3 Content validity

Each of the nodes in the model is necessary in order to make the burning decision. The boundaries of the states of each node were reviewed to ensure that there were no gaps or overlaps.

5.4 Concurrent validity

It is not possible to evaluate concurrent validity, as no other models (apart from [6]) were found with which one could compare the DBDN's behavior.

5.5 Convergent and discriminant validity

As mentioned in section 5.1, the structure of the DBDN is similar to the farmer example given in Clemen [19]. Within the DBDN, the three types of nodes making up "today's weather" are modelled in the same way. The two sets of nodes for "advance weather" are also modelled in the same way. The states of the nodes in the "delivery" section of the model are all discretized with the same accuracy of 0.25 day's cane quota.

5.6 Predictive validity

Behavior sensitivity tests. As Hugin only allows parameter sensitivity analyses to be performed on chance nodes, parameter sensitivity analyses were performed on the intermediate nodes connected to the decision. For example, the "sensitivity set" [13] for the node "Grower's belief of burning conditions, Day 1" shows that the previous day's actual rain type; the rain type forecast on the current day; and the grower's belief of cane dryness on the previous day play the largest role in the "Grower's belief of burning conditions" on Day 1. This makes sense, as rain plays a large role in the cane's dryness and the ability of the grower to burn. The results of the Value of Information analysis show that more information about the future weather conditions would improve the decision to burn. This is borne out by the fact that growers are always monitoring weather forecasts.

Extreme conditions tests. The following extreme conditions were identified:

- The model should never burn cane if the cane is damp or wet, if there is high wind, or if there is rain (drizzle or thundershowers).
- If sufficient cane is burned to meet the day's deliveries and/or future deliveries (if bad weather is expected), the model should also not burn additional cane.
- The model should never burn more than 4 day's delivery quota of cane.
- The model should not allow the grower to deliver more cane than is available.

Inspection of the CPTs of the decision and utility nodes showed that the burning would not occur in these first two conditions. It is impossible for the model to burn more than 4 days' cane, as this is the maximum state for the burning decision. These results, and tests of amount of cane available vs. amount of cane delivered, were borne out in programmatic tests of the model run using the Hugin API.

Qualitative features tests. During the focus group meeting, the DBDN model was run in Hugin with actual weather data from that day (rainy and cold), and expected weather data for the following days. For most of the nodes, the participants agreed that the model was behaving in the way that was expected. For the "Grower's belief of dryness of cane" node, they agreed that they would not be able to burn cane by the morning of the third day, as the model predicted. Knowing that there was more wet weather on the way a few days later, they thought that the cane would become dry enough to burn on the afternoon of the third day, and if they could burn cane in the afternoon, they would. However, the model only considers the burning decision once a day, in the morning. Considering a burning opportunity twice a day would mimic the growers' behavior better. The CPT values for the "Grower's belief of dryness of cane" and "Grower's belief of burning conditions" nodes could perhaps be adjusted so that the cane "dries slightly faster".

Results of a pre-run scenario based on October 2016 weather data were shown to the focus group (see Fig. 2[5]). This is Scenario 5 from [6]. However, when running this scenario, evidence was not entered for the "Grower's belief of dryness of cane" node, as was done in [6]. Instead, this node is treated as an unobservable variable. Its belief is inferred from the three previous day's actual weather nodes. An enhancement in version 3, compared to [6], is the ability to enter the number of days' cane quota needed in advance, based on longer term rain or wind forecasts (see third last row). The grower commented that this was a "perfect scenario" for that time of year; that the scenario showed very typical wet weather and wet cane, and how the growers struggle to keep on track with cane deliveries at that time of year. The scenario shows how the cane is only dry enough to burn on Day 2. On Day 3, it would have been dry enough to burn, except for the rule that cane cannot be burned on days when rain is expected. Because of Days 3-5's drizzle, the cane becomes increasingly wet. Day 6's lack of rain dried the cane to some extent, but it then became wetter due to drizzle and thundershowers. When cane was available (Days 3-6), it was delivered at the maximum delivery capacity of the grower.

[5] See Fig. 4 in Appendix C for a larger version.

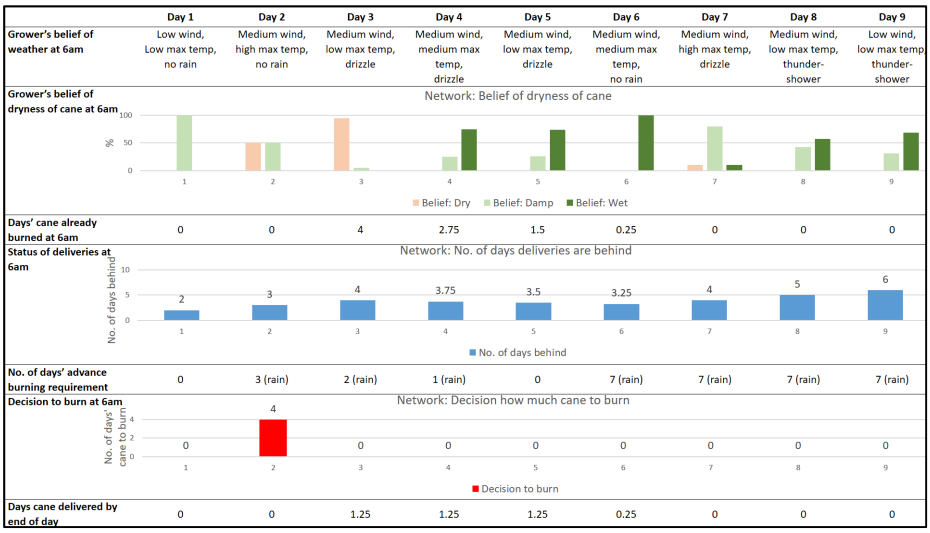

Fig. 2. Scenario based on weather data from October 2016 which was shown to the focus group participants. In this scenario, the grower starts 2 days behind with damp cane. Due to the rainy weather, the grower ends up 6 days behind. This was Scenario 5 in [6].[6]

5.7 "Write down" feedback from the focus group

In the "write down" after the focus group meeting, two participants said that the model assumptions were correct, whereas the other two mentioned that the large growers only burn every other day. Three of the five participants anonymously completed the Likert-scale question which asked to what extent the model represents what happens in reality in the mill's supply chain. Two participants scored it as "Well" whereas the other participant scored it as "Neutral". The fourth participant said that the model "represents 80% of what happens in reality." Responding to the question, "What aspects of the model need changing?", one participant said if possible, he would like to see the model attempting to burn cane in the afternoon if the cane had been too damp to burn in the morning. Another said, "More decision go into burning".

The participants then gave their comments about the model. One participant said, "Very good model. Have never seen something like this before. Thank you." Another said, "Excited to see it in action. Will be a great tool for growers and cane supply." Another participant was slightly more guarded, saying "It is a fair representation of what could happen." This sentiment was echoed by the fourth participant, who pointed out that with over 100 commercial growers delivering to the mill, each would have their own decision making criteria. However, "due to this, the model can only be 80% accurate".

[6] See Fig. 4 in Appendix C for a larger version.

6 Discussion and conclusion

With no other models with which to compare the DBDN, it was difficult to perform some of the validity tests in Pitchforth and Mengersen [9]'s validation framework. This problem was faced by other authors, e.g. [11, 12]. The framework provided a thorough set of tests with which to evaluate the model.

The pre-run scenario (Fig. 2) shows an improvement of version 3 of the model over version 1 [6] in that the number of days that the grower is behind is now correctly reflected, since the number of states for the status of deliveries has been extended from maximum '4 days behind' to '7 or more days behind'. However, the amount of cane burned in version 3 is less than what was burned in version 1. In version 1, one could burn on the same day as forecasted rain; and one could burn damp cane if one were behind with deliveries. In version 3, these assumptions were removed due to inputs from cane supply participants. The focus group participants noticed that the "Grower's belief of cane dryness" node in version 3 was not allowing burning as soon as they thought one could burn. It is easy to change the values of this node's CPTs.

Overall, the DBDN model shows a promising way of representing how the growers take the decision to burn cane, with two focus group participants saying this was "good", one saying it was "80%" and the other participant saying it was "neutral". However, this latter participant added that the model was a "fair representation of what could happen". The fact that these last two participants reported that more goes into the burning decision than currently modelled; and each of the mill's supply chain growers could have different decision criteria, shows that this is a complex problem. However, it is easy to add different factors to such BN models.

Suggestions for improving the model include adjusting the CPT values of the "Dryness of Cane" and "Grower's belief of burning conditions" nodes, so that it is more likely that the grower will burn cane. The large grower's burning rules can be adjusted to burning on alternate days (as is currently done for small growers); one can remove options to burn under one day's quota of cane. One could also consider allowing the grower to burn twice a day (morning and evening), as highlighted by one of the participants. This could be done by adding an additional decision node to each time slice. BNs' flexibility allows for all such changes to be made with relative ease. DBDNs offer a useful and apt way of modelling these kinds of adaptive recurring decisions. The model presented in this paper will provide a good starting point for representing the grower's daily burning decisions within a sugarcane supply chain simulation.

Acknowledgements

The authors gratefully thank: the cane procurement managers and focus group participants of the KwaZulu-Natal sugarcane supply chain for their inputs into and feedback about the model; University Capacity Development Programme (UCDP) for funding the upgrade to Hugin 8.6 and writing retreats; and College of Law and Management Studies, UKZN, for funding transport to the mill.

References

1. Smithers, J.: Review of sugarcane trash recovery systems for energy cogeneration in South Africa. Renew. Sust. Energy. Rev. 32, 915-925 (2014).
2. Sanjika, T.M., Bezuidenhout, C.N.: Driving factors-based approach for identifying performance indicators in sugarcane supply and processing systems. Br. Food J. 117, 1652-1669 (2015).
3. Kadwa, M., Bezuidenhout, C.N.: Modelling sugarcane supply consistency at a sugar mill. Comput. Electron. Agric. 111, 107-111 (2015).
4. Koller, D., Friedman, N.: Probabilistic graphical models: Principles and techniques. The MIT Press, Cambridge MA (2009).
5. Korb, K.B., Nicholson, A.E.: Bayesian artificial intelligence. CRC Press, Boca Raton FL (2011).
6. Price, C.S., Moodley, D., Pillay, A.W.: Dynamic Bayesian decision network to represent growers' adaptive pre-harvest burning decisions in a sugarcane supply chain. In: van Niekerk, J., Haskins, B. (eds.) SAICSIT '18 Conference, pp. 89-98. ACM, P.E., RSA (2018).
7. Drury, B., Rocha, C., Moura, M.-F., de Andrade Lopes, A.: The extraction from news stories: A causal topic centred Bayesian graph for sugarcane. IDEAS '16 Conference, pp. 364-369. ACM, Montreal QC, Canada (2016).
8. Marcot, B.G., Steventon, J.D., Sutherland, G.D., McCann, R.K.: Guidelines for developing and updating Bayesian belief networks applied to ecological modeling and conservation. Can. J. Forest Res. 36, 3063-3074 (2006).
9. Pitchforth, J., Mengersen, K.: A proposed validation framework for expert elicited Bayesian Networks. Expert Syst. Appl. 40, 162-167 (2013).
10. Pitchforth, J., Wu, P., Mengersen, K.: Applying a validation framework to a working airport terminal model. Expert Syst. Appl. 41, 4388-4400 (2014).
11. Goerlandt, F., Montewka, J.: A probabilistic model for accidental cargo oil outflow from product tankers in a ship–ship collision. Mar. Pollut. Bull. 79, 130-144 (2014).
12. Hänninen, M., Valdez Banda, O.A., Kujala, P.: Bayesian network model of maritime safety management. Expert Syst. Appl. 41, 7837-7846 (2014).
13. Coupé, V.M.H., Van Der Gaag, L.C.: Properties of sensitivity analysis of Bayesian belief networks. Ann. Math. Artif. Intel. 36, 323-356 (2002).
14. Price, C.S., Moodley, D., Bezuidenhout, C.N.: Using agent-based simulation to explore sugarcane supply chain transport complexities at a mill scale. In: Kruger, H.A. (ed.) 43rd Annual ORSSA Conference, pp. 88-96. Parys, RSA (2014).
15. Hugin: HUGIN Expert A/S (https://www.hugin.com/). (2019).
16. Hänninen, M.: Bayesian networks for maritime traffic accident prevention: Benefits and challenges. Accid. Anal. Prev. 73, 305-312 (2014).
17. NORSYS Software Corp.: Netica ™ Application (https://www.norsys.com/netica.html). (2019).
18. Greenbaum, T.L.: The handbook for focus group research. Sage, Thousand Oaks CA (1998).
19. Clemen, R.T.: Making hard decisions: An introduction to decision analysis. Duxbury, Belmont CA (1991).
20. Stray, B.J., van Vuuren, J.H., Bezuidenhout, C.N.: An optimisation-based seasonal sugarcane harvest scheduling decision support system for commercial growers in South Africa. Comput. Electron. Agric. 83, 21-31 (2012).
21. Le Masson, J.: Articulation de modèles de planification logistique et d'approvisionnement d'une sucrerie: Application à la mécanisation de la récolte de canne dans le bassin de

Noodsberg (Afrique du Sud). Masters dissertation, AgroParisTech (Paris) – Cirad (Montpellier), France (2007).
22. McDonald, B.C.: The development of a swarm intelligent simulation tool for sugarcane transport logistics systems. Masters dissertation, School of Computer Science, University of KwaZulu-Natal, Durban, South Africa (2008).

Appendix A: Focus group meeting methodology

The Cane Procurement Manager invited the participants from the mill's sugarcane supply chain. The five participants all had experience of growing cane, and most also had experience of cane supply from the mill's point of view as well (see Table 1). Apart from the Cane Procurement Manager, none of the participants had heard about the model before, and none had been interviewed by the first author before. The meeting was held at the mill in October 2019. On that day, the weather was cold and drizzly, as a cold front had just set in.

The focus group moderator (first author) gave all of the participants a set of documents. These contained the permission from the Mill Group Board to conduct research, the ethical clearance certification (HSS/0204/101), the letter requesting permission to participate and the consent form. A printout of the PowerPoint slides and an A3 copy of the DBDN (see Fig. 1) completed the document pack. After signing consent documents, the moderator introduced the concept of BNs. The DBDN model was introduced in a modular fashion, together with the assumptions that had led to the model's creation. The moderator took note of comments made and questions asked by the participants during and after the presentation. During this part of the focus group meeting, one of the participants had to leave.

Once the model and its assumptions had been introduced and discussed, the DBDN model was run in Hugin. The participants asked the moderator to enter that day's and the following days' data into the model to see how the model would react to cold, drizzly weather and wet cane. Thereafter the results of a scenario which had been run before the meeting were presented. The scenario was based on actual mill weather data from October 2016, since there had been much rain at that time; the cane was damp and the grower was behind with deliveries. This scenario is Scenario 5 of [6], and was chosen because spring rain interrupts deliveries.

At the conclusion of the meeting, the participants were asked to participate in a "write down" [18] so as to gauge each participant's view independently of others' views or of those dominating the discussion. In this exercise, they briefly answered questions on the model's assumptions, and what aspects of the model needed changing. They were asked to rate the model's representativeness of what happens in the mill's supply chain on a five point Likert scale question (Very poorly to Very well). Finally, they were asked for any additional comments on the model.

Appendix B: Model assumptions

The DBDN (version 3) has the following assumptions:

- Cane which is burned today can be delivered tomorrow.
- Large growers burn daily and deliver daily; small growers burn on alternate days and deliver daily.
- No more than 4 day's cane quota is ever burned at one time.
- Cane can be burned, harvested and delivered in multiples of 0.25 of a day's quota.
- Growers can deliver up to two day's cane quota on any one day.
- Growers deliver cane depending on how much cane has been burned (and harvested); the grower's maximum daily delivery rate, and the mill's call for cane.
- If the mill calls for 100% (normal) daily deliveries, or more, the grower may deliver more than his one day's quota, if cane is available (e.g. if the grower is behind with deliveries); however, if the mill calls for less than one day's cane, the grower will not deliver more than the amount called for.

Rain and wind affect the decision to burn in the following ways:

- Wet cane cannot be burned.
- Cane cannot be burned in high winds.
- Cane cannot be burned if the grower believes that drizzle or thundershowers are forecast for that day.
- However, if cane is dry, and the grower believes that no rain will be forecast, and that high winds are not forecast, cane can be burned.

Cane dryness is affected by the weather:

- Cane dries quickly after a thundershower, but soft soaking drizzle leaves the cane damp or wet.
- Wet or damp cane dries dependent on the maximum temperature and the amount of wind.

If wet or windy weather is forecast for the following week, the grower may burn more cane in advance so that there is cane to deliver when burning cannot take place. However, the grower will never burn more than four days' cane.

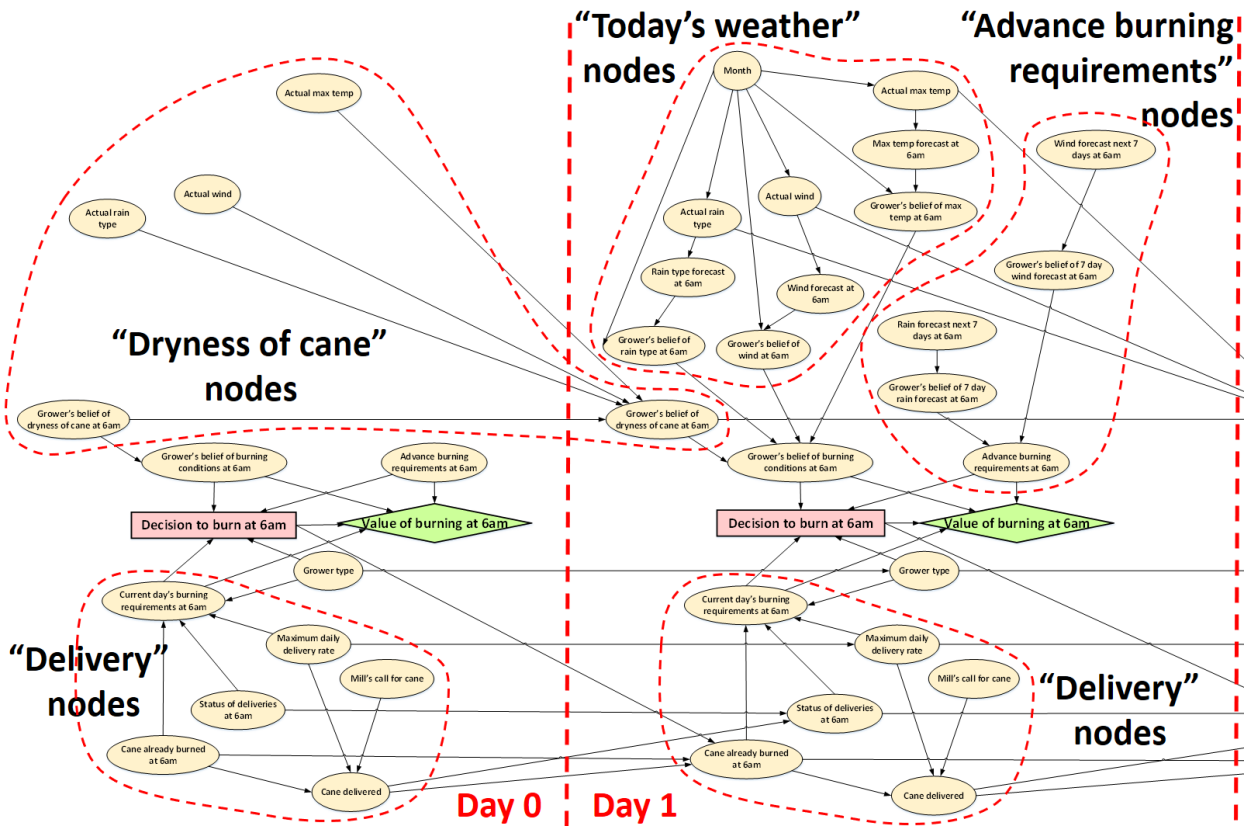

Fig. 3. (Larger version of Fig. 1). The dynamic Bayesian decision network which represents the grower's decision to burn cane before harvesting. Day 1's nodes are replicated for subsequent days.

Appendix C: Scenario

	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9
Grower's belief of weather at 6am	Low wind, Low max temp, no rain	Medium wind, high max temp, no rain	Medium wind, low max temp, drizzle	Medium wind, medium max temp, drizzle	Medium wind, low max temp, drizzle	Medium wind, medium max temp, no rain	Medium wind, high max temp, drizzle	Medium wind, low max temp, thunder-shower	Low wind, low max temp, thunder-shower
Grower's belief of dryness of cane at 6am	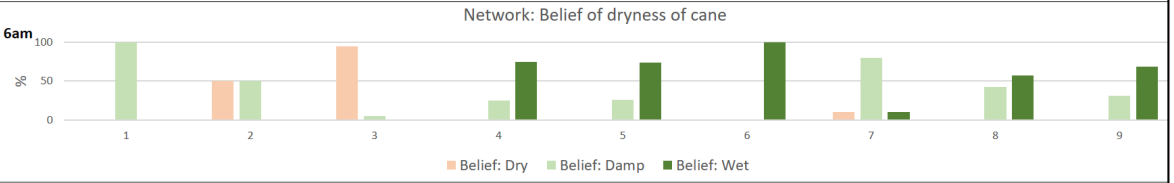								
Days' cane already burned at 6am	0	0	4	2.75	1.5	0.25	0	0	0
Status of deliveries at 6am	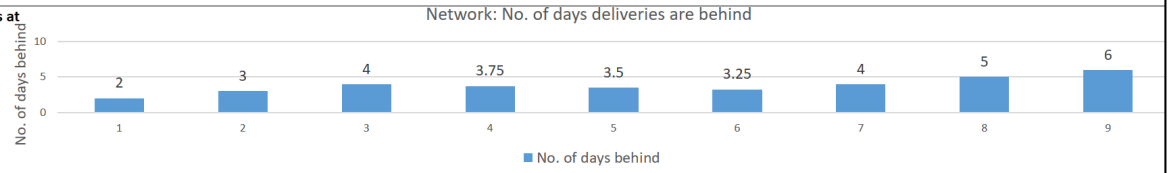								
No. of days' advance burning requirement	0	3 (rain)	2 (rain)	1 (rain)	0	7 (rain)	7 (rain)	7 (rain)	7 (rain)
Decision to burn at 6am	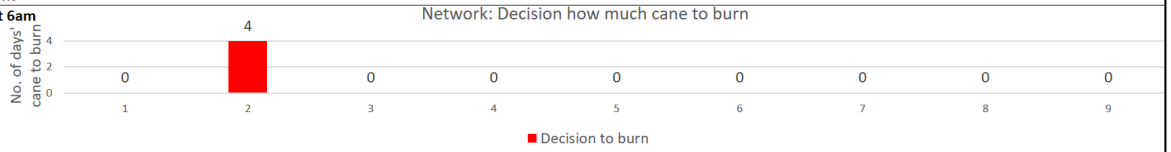								
Days cane delivered by end of day	0	0	1.25	1.25	1.25	0.25	0	0	0

Fig. 4. (Larger version of Fig. 2). Scenario based on weather data from October 2016 which was shown to the focus group participants. In this scenario, the grower starts 2 days behind with damp cane. Due to the rainy weather, the grower ends up 6 days behind. This was Scenario 5 in [6].