



# A Taxonomy of Explainable Bayesian Networks

Iena Petronella Derks<sup>1</sup>  and Alta de Waal<sup>1,2</sup> 

<sup>1</sup> Department of Statistics, University of Pretoria, Pretoria, South Africa  
inekederks1@gmail.com

<sup>2</sup> Center for Artificial Intelligence Research (CAIR), Pretoria, South Africa

**Abstract.** Artificial Intelligence (AI), and in particular, the explainability thereof, has gained phenomenal attention over the last few years. Whilst we usually do not question the decision-making process of these systems in situations where only the outcome is of interest, we do however pay close attention when these systems are applied in areas where the decisions directly influence the lives of humans. It is especially noisy and uncertain observations close to the decision boundary which results in predictions which cannot necessarily be explained that may foster mistrust among end-users. This drew attention to AI methods for which the outcomes can be explained. Bayesian networks are probabilistic graphical models that can be used as a tool to manage uncertainty. The probabilistic framework of a Bayesian network allows for explainability in the model, reasoning and evidence. The use of these methods is mostly ad hoc and not as well organised as explainability methods in the wider AI research field. As such, we introduce a taxonomy of explainability in Bayesian networks. We extend the existing categorisation of explainability in the model, reasoning or evidence to include explanation of decisions. The explanations obtained from the explainability methods are illustrated by means of a simple medical diagnostic scenario. The taxonomy introduced in this paper has the potential not only to encourage end-users to efficiently communicate outcomes obtained, but also support their understanding of how and, more importantly, why certain predictions were made.

**Keywords:** Bayesian network · Reasoning · Explainability

## 1 Introduction

Advances in technology have contributed to the generation of big data in nearly all fields of science, giving rise to new challenges with respect to explainability of models and techniques used to analyse such data. These models and techniques are often too complex; concealing the knowledge within the machine, hence decreasing the extent of interpretability of results. Subsequently, the lack of explainable models and techniques contribute to mistrust among users in fields of science where interpretability and explainability are indispensable.

To elucidate the need for explainable models, consider the following three scenarios. Firstly, suppose a medical diagnosis system is used to determine whether a tumour sample is malignant or benign. Here, the medical practitioner must be able to understand how and why the system reached the decision, and, if necessary, inspect whether the decision is supported by medical knowledge [2]. Next, consider self-driving cars. In this context, the self-driving car must be able to process information faster than a human, such that accidents and fatalities can be avoided [21]. Suppose a self-driving car is involved in an accident, then the system must be able to explain that in order to avoid hitting a pedestrian, the only option was to swerve out of the way and, by coincidence, into another vehicle. Lastly, consider an online restaurant review system, where reviews are classified as positive or negative based on the words contained in the review. Here, the classifier simply returns whether a review is positive or negative, without explaining which words contributed to the classification. As such, negative reviews that are expressed in, for example, a sarcastic manner, might be classified as positive, resulting in a restaurant receiving a higher rating and more diners – who might experience bad service (or even food poisoning) as a result of mislabelled reviews.

Given its relevance in many application areas, the explainability problem has attracted a great deal of attention in recent years, and as such, is an open research area [24]. The manifestation of explainable systems in high-risk areas has influenced the development of explainable artificial intelligence (XAI) in the sense of prescriptions or taxonomies of explanation. These include fairness, accountability, transparency and ethicality [3, 11, 23]. The foundation of such a system should include these prescriptions such that a level of usable intelligence is reached to not only understand model behaviour [1] but also understand the context of an application task [14]. Bayesian networks (BNs) – which lie at the intersection of AI, machine learning, and statistics – are probabilistic graphical models that can be used as a tool to manage uncertainty. These graphical models allow the user to reason about uncertainty in the problem domain by updating ones beliefs, whether this reasoning occurs from cause to effect, or from effect to cause. Reasoning in Bayesian networks is often referred to as what-if questions. The flexibility of a Bayesian network allows for these questions to be predictive, diagnostic and inter-causal. Some what-if questions might be intuitive to formulate, but this is not always the case especially on a diagnostic and inter-causal level. This might result in sub-optimal use of explainability in BNs – especially on an end-user level. Apart from well-established reasoning methods, the probabilistic framework of a Bayesian network also allows for explainability in evidence. These include most probable explanation and most relevant explanation. To extend on the existing explainability methods, we propose an additional approach which considers explanations concerned with the decision-base.

In this paper, we research the current state of explainable models in AI and machine learning tasks, where the domain of interest is BNs. In the current research, explanation is often done by principled approaches to finding explanations for models, reasoning, and evidence. Using this, we are able to formulate a taxonomy of explainable BNs. We extend this taxonomy to include explanation of decisions. This taxonomy will provide end-users with a set of tools to

better understand predictions made by BNs and will therefore encourage efficient communication between end-users. The paper is structured as follows. We first investigate the community and scope of explainability methods in Sect. 2. Thereafter, we introduce explanation in BNs, which includes the formulation of principled approaches, the theoretical properties associated therewith and a hands-on medical diagnosis example. Section 4 presents our newly formulated taxonomy of explainable BNs. The final section concludes the paper and includes a short discussion of future work.

## 2 Related Work

In application areas where erroneous decisions have a direct impact on livelihood, relying on systems where the predictions cannot be explained may not be an option. Explainability in such systems aids in establishing trust in not only circumstances where the system is used as a primary decision tool, but also cases where the system takes on a supportive role [28].

Over the past few years, explainability in AI systems has gained immense attention from the research community. This is reflected in the launch of various events and organisations. The Defense Advanced Research Projects Agency (DARPA) launched the Explainable Artificial Intelligence (XAI) initiative in 2016. The XAI programs intention is to encourage the production of AI techniques where emphasis is placed on developing more accurate and precise models, while still maintaining a high level of explainability. Ultimately, XAI systems must be able to explain their rationale and enable understanding [12]. Conferences, such as the International Joint Conferences on Artificial Intelligence (IJCAI) conducts workshops specifically focusing on XAI [26]. This topic has also made a noticeable appearance at the Neural Information Processing Systems (NeurIPS) conference, with panel discussions solely concentrating on XAI.

The scope of explainability is inherently linked to the complexity of the model, as well as the goal thereof. Usually, but not necessarily, there is a trade-off between model accuracy and explainability – the higher the accuracy, the lower the explainability [31]. For example, decision trees provide a clear explanation but are often less accurate than deep learning models, which are less transparent. It should be mentioned that this trade-off is also connected to the quality of data. AI and machine learning models that are transparent by design, such as linear regression, decision trees and k-nearest neighbours, convey a degree of explainability [1]. However, when AI and machine learning models do not provide clear explanations, separate explainability methods are applied to the model to gain meaningful explanations. Methods of explainability are not limited to the behaviour of the model or decision-making process as a whole, and may be applied to single instances, predictions or decisions [6]. These explanations can be in the form of visualisations or natural language [10]. Some of the existing explainability methods are layer-wise relevance propagation (LRP), which are often used in deep neural networks where the prediction made by the network is propagated back into the neural network using a set of

predefined rules [27]. Another explainability method is local interpretable model-agnostic explanations (LIME). LIME methods can be used to explain prediction instances by attempting to understand the behaviour of the prediction function in the context of the prediction. Here, the user is able to obtain a local explanation for that particular instance. LIME can also be used to obtain explanations for the entire model by generating multiple instances [17]. Methods of explainability are also extended to document classifiers, where documents are classified based on predicted likelihood. Here, explanations can be produced based on a search through the text-space of possible word combinations – starting with a single word and expanding the number of words until an explanation is found [25].

Uncertainty is present in the majority of AI fields, such as knowledge representation, learning and reasoning [22]. Real-world data often contain noisy and uncertain observations close to the decision boundary, which may result in predictions that cannot be explained [7]. Probabilistic graphical models can be seen as uncertainty management tools as they are able to represent and reason with uncertainty. These probabilistic models are often employed to support decision making in various application fields, including legal and medical applications [29]. One such probabilistic model is BNs, which is capable of combining expert knowledge and statistical data, therefore allowing for complex scenarios to be modelled. However, not only are the inner workings of Bayesian networks complicated to most end-users [15], the explanation of probabilistic reasoning is challenging and as such results appear to be counter-intuitive or wrong [16]. Therefore, there exists a demand for explanation in Bayesian networks.

Explanation methods for Bayesian networks can be divided into three broad approaches. The first approach consists of presenting information contained in the knowledge base and is known as *explanation of the model*. There are two objectives associated with this type of explanation. Firstly, explanation of the model is used to assist application experts in the model-construction phase. Secondly, it is used for instructional purposes to offer knowledge about the domain [19]. The objective of the second approach is to justify the conclusion and how it was obtained. This approach is referred to as *explanation of reasoning* [9]. The final approach, *explanation of evidence*, is concerned with the treatment of the variables in the Bayesian network [13]. In explanation of evidence, also referred to as *abduction*, an explanation is seen as the configuration of a portion of the variables present in the Bayesian network, given evidence. Not included in the aforementioned explanation methods are techniques that describe whether the end-user is ready to make a decision, and if not, what additional information is required to better prepare for decision making. Techniques such as sensitivity analysis [4] and same-decision probability [5] provide the end-user with insight on decisions. We group these methods into a fourth approach, *explanation of decisions*. For the purpose of this paper, and accordingly, the formulation of the explainable taxonomy, we only consider explanation of reasoning, evidence, and decisions. Explanation of the model is excluded from this taxonomy – at the time being – as the intent of the taxonomy is to support understanding of how and why predictions were made and not on the model-construction itself.

### 3 Explainable Bayesian Networks

We adapt the XAI terminology to the scope of BNs by defining the term XBN and thereby referring to explainable BNs. To illustrate XBN in Bayesian networks, consider the Asia network from Lauritzen and Spiegelhalter (1988) [20] as an example.

**Example Statement:** Suppose a patient visits a doctor, complaining about shortness of breath (dyspnoea) (**D**). The patient is worried he might have lung cancer. The doctor knows that lung cancer is only one of the possible causes for dyspnoea, and other causes include bronchitis (**B**) and tuberculosis (**T**). From her training, the doctor knows that smoking increases the probability of lung cancer (**C**) and bronchitis. Both tuberculosis and lung cancer would result in an abnormal X-ray (**X**) result. Lastly, a recent visit to Asia might increase the probability of tuberculosis, as the disease is more prevalent there than the patient’s country of origin.

From this example statement, the nodes and values are defined and then the graphical structure of the BN is constructed. This is followed by the quantification of the conditional probability tables (CPTs) for each node [18]<sup>1</sup>. The final BN is illustrated in Fig. 1. Now that the domain and uncertainty are represented in the BN, we will look into how to use the BN. Reasoning in BNs takes place once we observe the value of one or more variables and we want to condition on this new information [18]. It is important to note that this information need not necessarily flow in the direction of the arcs, and therefore, reasoning can occur in the opposite direction of the arcs.




Fig. 1. Asia Bayesian network

<sup>1</sup> All BN models are constructed in BayesiaLab ([www.bayesia.com](http://www.bayesia.com)).

### 3.1 Explanation of Reasoning

Suppose during the doctor’s appointment, the patient tells the doctor he is a smoker before any symptoms are assessed. As mentioned earlier, the doctor knows smoking increases the probability of the patient having lung cancer and bronchitis. This will, in turn, also influence the expectation of other symptoms, such as the result of the chest X-Ray and shortness of breath. Here, our reasoning is performed from new information about the causes to new beliefs of the effects. This type of reasoning is referred to as **predictive reasoning** and follows the direction of the arcs in the network. Through predictive reasoning, we are interested in questions concerning *what will happen*. In some cases, predictive reasoning is not of great insight and it is often required to reason from symptoms (effect) to cause, which entails information flow in the opposite direction to the network arcs. For example, bronchitis can be seen as an effect of smoking. Accordingly, we are interested in computing  $P(S|B)$ . This is referred to as **diagnostic reasoning** and is typically used in situations where we want to determine *what went wrong*. The final type of probabilistic reasoning in BNs is **inter-causal reasoning**, which relates to mutual causes of a common effect – typically indicated by a v-structure in the network. In other words, inference is performed on the parent nodes of a shared child node. Note that the parent nodes are independent of one another unless the shared child node is observed, a concept known as *d-separation*. From the Asia network, we observe a v-structure between Tuberculosis, Lung Cancer and Tuberculosis or Cancer (see Fig. 2a). Here, Tuberculosis is independent from Lung cancer. Suppose we observe the patient has either Tuberculosis or Cancer – indicated by the green (or light grey if viewed in grey-scale) bar in Fig. 2b – then this observation increases the probabilities of the parent nodes, Tuberculosis and Lung Cancer. However, if it is then revealed that the patient does, in fact, have Tuberculosis it, in turn, lowers the probability of a patient having Lung Cancer (see Fig. 2c). We can then say Lung Cancer has been *explained away*. It should be noted that the probabilistic reasoning methods discussed above can be used as is, or can be combined to accommodate the problem at hand.

### 3.2 Explanation of Evidence

Sometimes, users of the system find the results of reasoning unclear or questionable. One way to address this is to provide scenarios for which the reasoning outcomes are upheld. A fully specified scenario is easier to understand than a set of reasoning outcomes. Explanation of evidence methods are useful in specifying these scenarios. They are based on the posterior probability and the generalised Bayes factor. Firstly, we focus on methods that aim to find a configuration of variables such that the posterior probability is maximised given the evidence. Here, we consider the Most Probable Explanation (MPE), which is a special case of the Maximum A Posteriori (MAP). The MAP in a BN is a variable configuration which includes a subset of unobserved variables in the explanation set

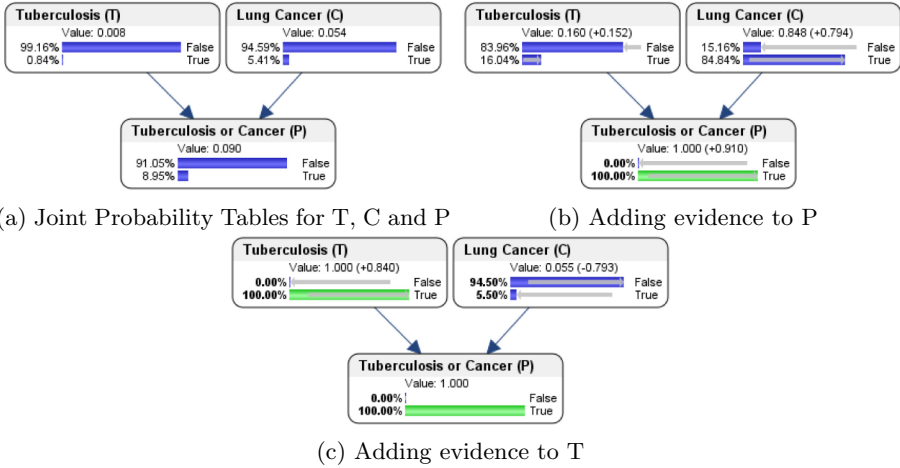


Fig. 2. Belief updating for T, C and P

such that the posterior probability – given evidence – is maximised. Similarly, if the variable configuration consists of all variables present in the explanation set, we have an MPE solution [13]. However, in some real-world applications, the variable set often consists of a large number of variables, which may result in over-specified or under-specified explanations obtained from the MPE. In fact, only a few variables may be relevant in explaining the evidence. The next approach finds a single instantiation that maximises the generalised Bayes factor in a trans-dimensional space containing all possible partial instantiations. In other words, this approach aims to obtain an explanation only consisting of the most relevant variables in the BN, given the evidence. This approach is known as the Most Relevant Explanation (MRE) [32–34].

**Most Probable Explanation.** Let’s first consider the MPE method. Recall that the MPE finds the complete instantiation of the target variables – which are defined to be unobserved – such that the joint posterior probability is maximised given evidence. Figure 3 shows the scenario (or case) that has the highest joint probability in the Asia network. Note here the probabilities are replaced by the likelihood of the variable state belonging to the most probable scenario, for example, if we look at the two possible states for Bronchitis, we see that ‘False’, i.e., the patient does not have bronchitis, is more probable. Suppose we discover the patient suffers from shortness of breath, we can then set the evidence for Dyspnoea as ‘True’ (illustrated in Fig. 4). By introducing this new evidence, we now observe a slightly different scenario, where it is more probable for the patient to be a smoker and have bronchitis. Notice here that variables that seem irrelevant to the evidence explanation, such as Visit to Asia and XRay, are included in the explanation. This could lead to overspecified hypotheses, especially in larger networks.

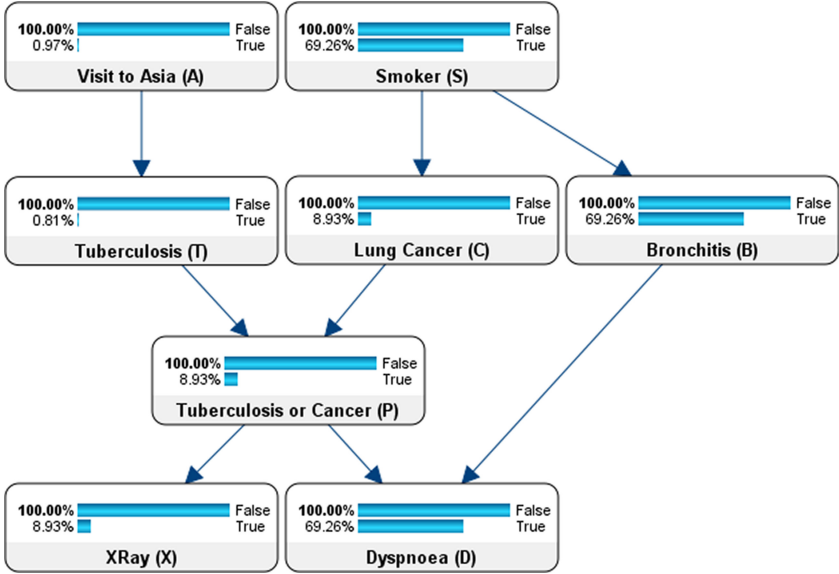


Fig. 3. Initial MPE for Asia Bayesian network

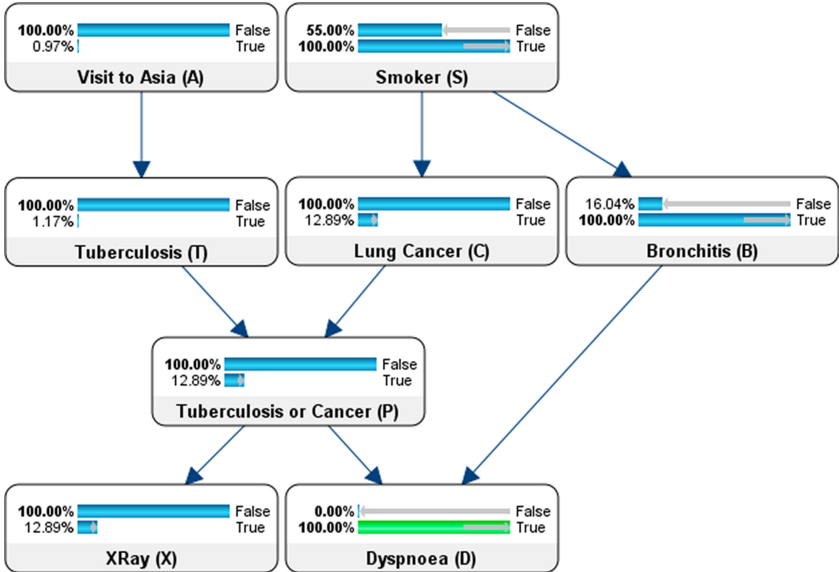


Fig. 4. Updated MPE for Asia Bayesian network

**Most Relevant Explanation.** To avoid an overspecified hypotheses, one approach is to trim or prune less relevant variables from the explanation. That is, instead of finding the complete instantiation of the target variables, a partial



instantiation of the target variables is found such that the *generalised Bayes factor* is maximised. Let's first consider the explanations obtained from the generalised Bayes factor. Again, suppose the patient suffers from shortness of breath (evidence). We are then interested in finding only those variables that are relevant in explaining why the patient has shortness of breath. Table 1 contains the set of explanations obtained from the generalised Bayes factor. For example, the last entry shows that a possible explanation for shortness of breath is a trip to Asia and an abnormal X-ray. Thus including only 2 variables from the remaining 7 variables (excluding Dyspnoea). As mentioned, the MRE is the explanation that maximises the generalised Bayes factor. From Table 1 we see that having Bronchitis best explains the shortness of breath. Notice that this explanation does not include Smoking, as opposed to the MPE which included Smoking. Thus, although smoking is a probable cause for shortness of breath, it is not the most relevant cause. An interesting characteristic of the MRE is its ability to capture the *explaining away* phenomenon [33].

**Table 1.** Explanations of GBF scores for Asia network

Explanation	Generalised Bayes factor
<b>(Bronchitis)</b>	<b>6.1391</b>
(Smoker, Tuberculosis or Cancer)	1.9818
(Tuberculosis or Cancer)	1.9771
(Lung Cancer, Smoker)	1.9723
(Lung Cancer)	1.9678
(Smoker, Tuberculosis)	1.8896
(Tuberculosis)	1.8276
(Smoker, XRay)	1.7779
(Smoker)	1.7322
(Visit to Asia, XRay)	1.5635

### 3.3 Explanation of Decisions

Hidden or unobserved variables appear in most application fields, especially in areas where decisions made by the end-user directly influence human lives. For example, when first examining a patient, the health-state of the patient is unknown. In these situations, one would typically ask two questions. The first being *given the available information, are we ready to make a decision?* and secondly, *if we are not yet ready to make a decision, what additional information do we require to make an informed decision?*. To answer these questions, the authors of [5] propose a threshold-based notion, named *same-decision probability*, which provides the user with a confidence measure that represents the probability that a certain decision will be made, had information pertaining unknown variables

been made available. Another possible threshold-based solution to this is sensitivity analysis [30]. In sensitivity analysis, the assessments for the conditional probabilities in the BN are systematically changed to study the effect on the output produced by the network. The idea is that some conditional probabilities will hardly influence the decisions, while others will have significant impact.

**Same-Decision Probability.** Suppose we are interested in making a decision on whether the patient is a smoker (Smoking), which is conditioned on evidence Tuberculosis or Cancer. We can then use the BN such that our decision pertaining to the hypothesis is supported on the basis that the belief in the hypothesis given some evidence exceeds a given threshold. Now, the patient may have access to information that is unknown to us, for example, the patient recently visited Asia and chose not to disclose this information. Therefore, we do not have access to the true state of this variable. The true state knowledge may confirm or contradict our decision based on the probability of smoking given some evidence and the patient visiting Asia. If we now compare this probability with some threshold, we have a degree of confidence in our original decision regarding smoking and the available evidence. Had the patient disclosed his trip to Asia, it is then unlikely that we would have made a different decision. Hence, we can make use of the same-decision probability (SDP). Consider now the BN given in Fig. 5. Notice here the addition of three nodes,  $P(\text{Smoker} = \text{True})$ ,  $\text{Decision Threshold}$  and  $\text{Decision}$ . Where  $P(\text{Smoker} = \text{True})$  represents the hypothesis probability and the decision threshold is set to 55%. Suppose now we update our network such that Tuberculosis or Cancer (P) is True – to reflect the scenario discussed above. The hypothesis probability then increases from 50.00% to 84.35% (see Fig. 6). Our decision is confirmed given the threshold since the hypothesis probability now exceeds the given threshold value. From Table 2, the SDP before adding evidence for the ‘True’ state is 0.00%. After adding evidence, the SDP for our decision is 83.88%, indicating that our decision confidence is 83.88%<sup>2</sup>.

**Table 2.** Decision confidence for Asia network

	States	Minimum	Maximum	Mean	Standard deviation
No evidence	False	100.00%	100.00%	100.00%	0.00%
	True	0.00%	0.00%	0.00%	0.00%
Evidence	False	0.00%	100.00%	16.12%	36.77%
	True	0.00%	100.00%	83.88%	36.77%

<sup>2</sup> The SDP scenario was constructed using the decision node functionality in Bayesialab. The decision nodes are indicated as green (or dark grey if viewed in grey-scale).

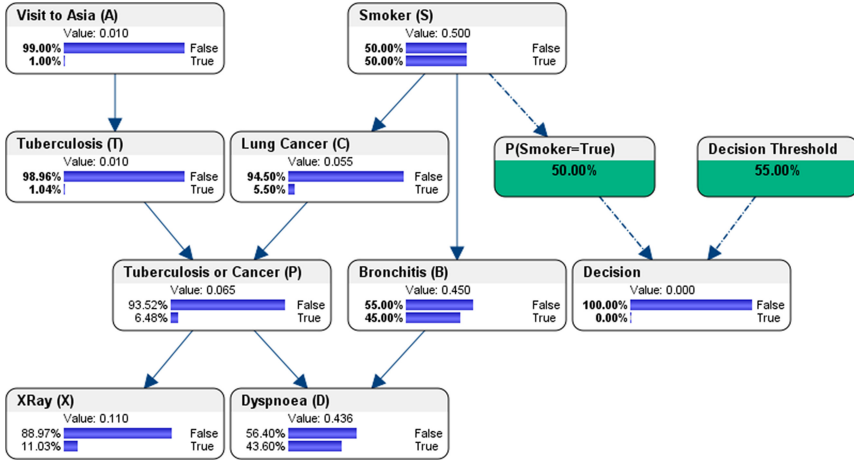


Fig. 5. Addition of decision node in Asia network

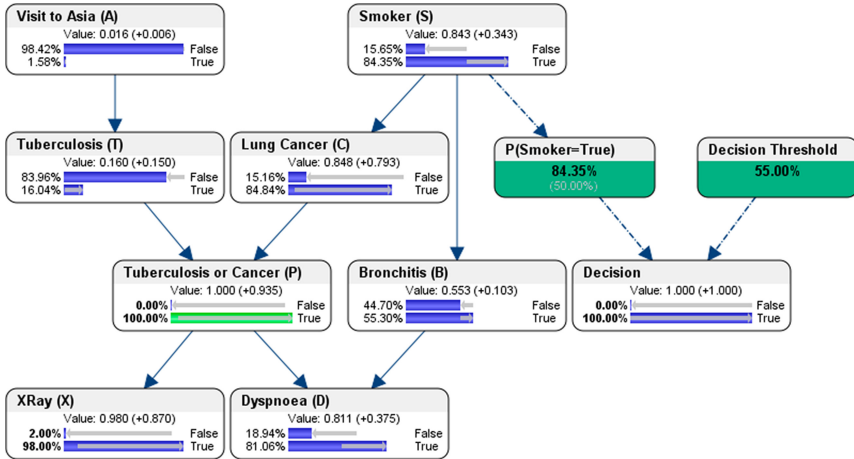


Fig. 6. Updated decision for Asia network

## 4 XBN in Action

The point of XBN is to explain the AI task at hand. In other words, the question the decision-maker seeks to answer, and not the technique in principle. Therefore, we need to be able to freely ask ‘why’ or ‘what’ and from this select a method that would best address the AI task. In Fig. 7 we present a taxonomy of XBN. The purpose of this taxonomy is to categorise XBN methods into four phases of BNs: The first phase involves the construction of the BN model. Explanation in the ‘model’ phase is critical when the model is based on expert knowledge. The second phase is reasoning, the third phase evidence, and the

fourth decision. Explanation of the model and sensitivity analysis are illustrated in grey as it is out of scope for this paper. Although we define the taxonomy along these phases, we do acknowledge that not all phases are necessarily utilised by the decision-maker. For example, when using BNs to facilitate participatory modelling [8], the main emphasis is on explaining the model. Or, when using BNs as a classifier, the emphasis is on explaining the decisions. In this section, we present typical questions of interest to the decision-maker in each category of the XBN taxonomy.

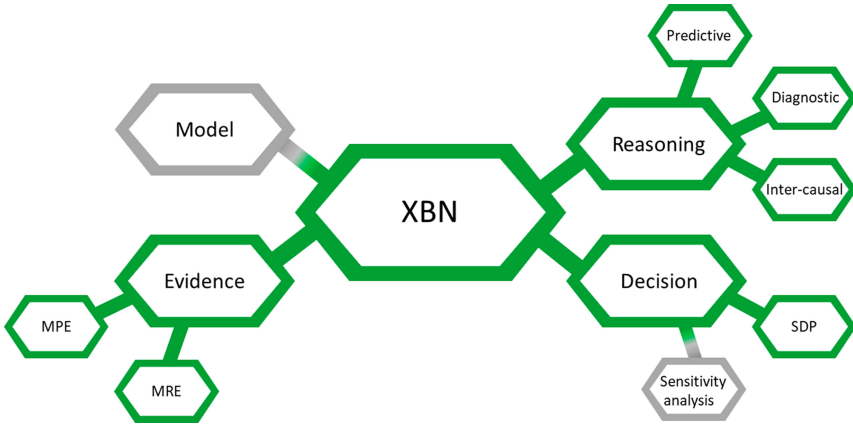


Fig. 7. A schematic view of XBN

#### 4.1 Reasoning

Reasoning in the XBN taxonomy is concerned with the justification of a conclusion. Returning to our Asia example, the end-user might ask the following question,

*“Given the patient recently visited Asia, how likely is an abnormal chest X-Ray?”*

Here, we are concerned with a single outcome: the X-Ray result. On the other hand, the doctor may have knowledge about symptoms presented by the patient and ask,

*“What is the probability of a patient being a smoker, given that he presented shortness of breath?”*

We can extend this to a forensic context. Suppose a crime scene is investigated where a severely burned body is found. The forensic analyst can then ask,

*“The burn victim is found with a protruded tongue, was the victim exposed to fire before death or after?”*

Consider now a financial service context where a young prospective home owner is declined a loan. The service provider can then ask,

*“Did the prospective owner not qualify for the home loan because of his age?”*

From these examples, we see that explanation of reasoning is used where questions are asked in the context of single variable outcomes for diagnosis.

## 4.2 Evidence

When we are interested in the subset of variables that describes specific scenarios, we use explanation of evidence methods. For example, in our Asia example the doctor may ask,

*“Which diseases are most probable to the symptoms presented by the patient?”*

or

*“Which diseases are most relevant to the symptoms presented by the patient?”*

In a forensic context, the forensic analyst investigating a crime scene may ask the following question,

*“What are the most relevant causes of death, given the victim is found with a severely burned body and protruded tongue?”*

Similarly this can be applied to fraud detection. Suppose the analyst investigates the credit card transactions of a consumer. The analyst can then ask,

*“What are the most probable transaction features that contributed to the flagging of this consumer?”*

Explanation of evidence can also be used to provide explanations for financial service circumstances. For example, if a prospective home owner is turned down for a loan, he may ask the service provider which features in his risk profile are more relevant (contributed most) to being turned down.

## 4.3 Decisions

Explanation of decisions typically asks the following questions *“Do we have enough evidence to make a decision?”*, and if not, *“what additional evidence is required to make a decision?”*. For example, in our Asia example we can ask,

*“Do we have enough evidence on the symptoms presented to make a decision on the disease?”*

or

*“Since we are not yet able to determine the disease, what additional information – test, underlying symptoms, comorbidities – is required to make a decision?”*

Applied to forensic investigations, this can be used to answer questions relating to crime scene investigations. The analyst may ask questions regarding the actual evidence collected from the crime scene, i.e., if enough evidence is collected to rule a crime as a homicide or what additional evidence is required to rule the crime as a homicide. Should they investigate further or is the evidence that is already collected enough to make an informed decision?

## 5 Conclusion

The development of AI systems has faced incredible advances in recent years. We are now exposed to these systems on a daily basis, such as product recommendation systems used by online retailers. However, these systems are also being implemented by medical practitioners, forensic analysts and financial services – application areas where decisions directly influence the lives of humans. It is because of these high-risk application areas that progressively more interest is given to the explainability of these systems.

This paper addresses the problem of explainability in BNs. We first explored the state of explainable AI and in particular BNs, which serves as a foundation for our XBN framework. We then presented a taxonomy to categorise XBN methods in order to emphasise the benefits of each method given a specific usage of the BN model. This XBN taxonomy will serve as a guideline, which will enable end-users to understand how and why predictions were made and will, therefore, be able to better communicate how outcomes were obtained based on these predictions.

The XBN taxonomy consists of explanation of reasoning, evidence and decisions. Explanation of the model is reserved for future work, since the taxonomy described in this paper is focused on how and why predictions were made and not on the model-construction phase. Other future research endeavours include the addition of more dimensions and methods to the XBN taxonomy – this involves more statistical-based methods and the incorporation of causability (which also addresses the quality of explanations) – as well as applying this taxonomy to real-world applications.

## References

1. Barredo Arrieta, A., et al.: Explainable artificial Intelligence (XAI): concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **58**, 82–115 (2020). <https://doi.org/10.1016/j.inffus.2019.12.012>
2. Brito-Sarracino, T., dos Santos, M.R., Antunes, E.F., de Andrade Santos, I.B., Kasmanas, J.C., de Leon Ferreira, A.C.P., et al.: Explainable machine learning for breast cancer diagnosis. In: 2019 8th Brazilian Conference on Intelligent Systems (BRACIS), pp. 681–686. IEEE (2019)
3. Cath, C.: Governing artificial intelligence: ethical, legal and technical opportunities and challenges. *Philos. Trans. R. Soc. Math. Phys. Eng. Sci.* **376**(2133) (2018). <https://doi.org/10.1098/rsta.2018.0080>

4. Chan, H., Darwiche, A.: On the robustness of most probable explanations. In: Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence, UAI 2006, June 2012
5. Choi, A., Xue, Y., Darwiche, A.: Same-decision probability: a confidence measure for threshold-based decisions. *Int. J. Approximate Reasoning* **53**(9), 1415–1428 (2012)
6. Das, A., Rad, P.: Opportunities and Challenges in Explainable Artificial Intelligence (XAI): A Survey. arXiv preprint [arXiv:2006.11371](https://arxiv.org/abs/2006.11371) (2020)
7. De Waal, A., Steyn, C.: Uncertainty measurements in neural network predictions for classification tasks. In: 2020 IEEE 23rd International Conference on Information Fusion (FUSION), pp. 1–7. IEEE (2020)
8. Düsöhl, M., Frank, S., Döll, P.: A review of Bayesian networks as a participatory modeling approach in support of sustainable environmental management. *J. Sustain. Dev.* **5**(12), 1 (2012). <https://doi.org/10.5539/jsd.v5n12p1>
9. Gallego, M.J.F.: Bayesian networks inference: Advanced algorithms for triangulation and partial abduction (2005)
10. Goebel, R., et al.: Explainable AI: the new 42? In: Holzinger, A., Kieseberg, P., Tjoa, A.M., Weippl, E. (eds.) CD-MAKE 2018. LNCS, vol. 11015, pp. 295–303. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-99740-7\\_21](https://doi.org/10.1007/978-3-319-99740-7_21)
11. Greene, D., Hoffmann, A.L., Stark, L.: Better, nicer, clearer, fairer: a critical assessment of the movement for ethical artificial intelligence and machine learning. In: Proceedings of the 52nd Hawaii International Conference on System Sciences, pp. 2122–2131 (2019). <https://doi.org/10.24251/hicss.2019.258>
12. Gunning, D., Aha, D.W.: DARPA’s explainable artificial intelligence program. *AI Mag.* **40**(2), 44–58 (2019)
13. Helldin, T., Riveiro, M.: Explanation methods for Bayesian networks: review and application to a maritime scenario. In: Proceedings of The 3rd Annual Skövde Workshop on Information Fusion Topics, SWIFT, pp. 11–16 (2009)
14. Holzinger, A., et al.: Towards the Augmented Pathologist: Challenges of Explainable-AI in Digital Pathology. arXiv preprint [arXiv:1712.06657](https://arxiv.org/abs/1712.06657) pp. 1–34 (2017). <http://arxiv.org/abs/1712.06657>
15. Keppens, J.: Explaining Bayesian belief revision for legal applications. In: JURIX, pp. 63–72 (2016)
16. Keppens, J.: Explainable Bayesian network query results via natural language generation systems. In: Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law, pp. 42–51 (2019)
17. Khedkar, S., Subramanian, V., Shinde, G., Gandhi, P.: Explainable AI in healthcare. In: Healthcare (April 8, 2019). 2nd International Conference on Advances in Science and Technology (ICAST) (2019)
18. Korb, K.B., Nicholson, A.E.: Bayesian Artificial Intelligence. CRC Press, Boca Raton (2010)
19. Lacave, C., Díez, F.J.: A review of explanation methods for Bayesian networks. *Knowl. Eng. Rev.* **17**(2), 107–127 (2002). <https://doi.org/10.1017/S026988890200019X>
20. Lauritzen, S.L., Spiegelhalter, D.J.: Local computations with probabilities on graphical structures and their application to expert systems. *J. R. Stat. Soc. Ser. B (Methodological)* **50**(2), 157–194 (1988)
21. Lawless, W.F., Mittu, R., Sofge, D., Hiatt, L.: Artificial intelligence, autonomy, and human-machine teams: interdependence, context, and explainable AI. *AI Mag.* **40**(3), 5–13 (2019)

22. Lecue, F.: On the role of knowledge graphs in explainable AI. *Seman. Web* **11**(1), 41–51 (2020)
23. Leslie, D.: Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector (2019). <https://doi.org/10.5281/zenodo.3240529>
24. Lipton, Z.C.: The mythos of model interpretability. *Queue* **16**(3), 31–57 (2018)
25. Martens, D., Provost, F.: Explaining data-driven document classifications. *MIS Q.* **38**(1), 73–100 (2014)
26. Miller, T., Weber, R., Magazzeni, D.: Proceedings of the IJCAI 2019 Workshop on Explainable AI (2019)
27. Montavon, G., Binder, A., Lapuschkin, S., Samek, W., Müller, K.-R.: Layer-wise relevance propagation: an overview. In: Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Müller, K.-R. (eds.) *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. LNCS (LNAI), vol. 11700, pp. 193–209. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-28954-6\\_10](https://doi.org/10.1007/978-3-030-28954-6_10)
28. Samek, W., Müller, K.-R.: Towards explainable artificial intelligence. In: Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Müller, K.-R. (eds.) *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. LNCS (LNAI), vol. 11700, pp. 5–22. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-28954-6\\_1](https://doi.org/10.1007/978-3-030-28954-6_1)
29. Timmer, S.T., Meyer, J.J.C., Prakken, H., Renooij, S., Verheij, B.: A two-phase method for extracting explanatory arguments from Bayesian networks. *International Journal of Approximate Reasoning* **80**, 475–494 (2017)
30. van der Gaag, L.C., Coupé, V.M.H.: Sensitivity analysis for threshold decision making with Bayesian belief networks. In: Lamma, E., Mello, P. (eds.) *AI\*IA 1999*. LNCS (LNAI), vol. 1792, pp. 37–48. Springer, Heidelberg (2000). [https://doi.org/10.1007/3-540-46238-4\\_4](https://doi.org/10.1007/3-540-46238-4_4)
31. Xu, F., Uszkoreit, H., Du, Y., Fan, W., Zhao, D., Zhu, J.: Explainable AI: a brief survey on history, research areas, approaches and challenges. In: Tang, J., Kan, M.-Y., Zhao, D., Li, S., Zan, H. (eds.) *NLPCC 2019, Part II*. LNCS (LNAI), vol. 11839, pp. 563–574. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-32236-6\\_51](https://doi.org/10.1007/978-3-030-32236-6_51)
32. Yuan, C.: Some properties of most relevant explanation. In: *ExaCt*, pp. 118–126 (2009)
33. Yuan, C., Lim, H., Lu, T.C.: Most relevant explanation in Bayesian networks. *J. Artif. Intell. Res.* **42**, 309–352 (2011). <https://doi.org/10.1613/jair.3301>
34. Yuan, C., Liu, X., Lu, T.C., Lim, H.: Most relevant explanation: Properties, algorithms, and evaluations. In: *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence, UAI 2009*, pp. 631–638 (2009)