Subjective and temporal quality-of-life information in Bayesian networks

Johan Kwisthout[1][0000−0003−4383−7786]

Radboud University, PO Box 9104, 6500HE Nijmegen, The Netherlands
johan.kwisthout@donders.ru.nl
https://www.socsci.ru.nl/johank/index.html

Abstract. In clinical decision support systems, Bayesian networks are often preferred as formalism for learning and reasoning with causally related information. As white-box AI models they combine domain expertise and data in stochastic causal models that allow for transparent, robust, and explainable inferences, like the best fitting treatment or follow-up of cancer patients. To facilitate personalized healthcare, the models require both medical and quality-of-life aspects. A major challenge with quality-of-life information is that it is inherently subjective and ordinal; we are typically interested in change rather than absolute value; e.g., a treatment that slows down a decreasing quality-of-life. One needs to reason with subjective, but ordered and temporal information ("I'm feeling better today than yesterday"). This implies novel representation, learning, inferring, and explanation of such information in Bayesian networks. In this short paper I present some preliminary ideas on how to represent such information and show, using simple examples, how models can be trained on heterogeneous subjective data and how inference queries can be computed.

Keywords: Bayesian Networks · Quality-of-life · Decision Support Systems · Temporal Inferences · Subjective Information.

1 Introduction

Allowing people to live longer in good health, defined by their own quality-of-life (QoL) expectations, is a crucial mission within the Dutch Knowledge and Innovation Agenda on healthcare. Personalised treatment and follow-up is vital for both optimising the patient's well-being as well as mitigating the costs of healthcare [7]. Artificial intelligence, in the form of clinical decision support systems (CDSSs), can help healthcare practitioners and patients to make informed decisions about personalised care [18]. In these CDSSs, Bayesian networks are often preferred as formalism for learning and reasoning with causally related information [9]. As white-box AI models they combine domain expertise and data in stochastic causal models that allow for transparent, robust, and explainable inferences, like the best fitting treatment or follow-up of cancer patients. The ENDORISK model [15], for example, allows oncologists to estimate the risk

of preoperative lymph node metastasis; this allows for risk stratification and the avoidance of unnecessary removal of the lymph nodes in low-risk patients, considerably affecting their post-operative health-related QoL.

To facilitate decision making about personalized healthcare, the Bayesian networks at the heart of these CDSSs need to integrate both medical and QoL aspects [21]. Revicki and colleagues define QoL as *"a broad range of human experiences related to one's overall well-being. It implies value based on subjective functioning in comparison with personal expectations and is defined by subjective experiences, states and perceptions. Quality of life, by its very natures, is idiosyncratic to the individual, but intuitively meaningful and understandable to most people"* [16, p. 888]. A major challenge with QoL information, such as data obtained with patient questionnaires, is that it is inherently subjective and relative, as the above definition illustrates. For example, in the EORTC QoL questionnaire[1], participants may report on pain in the last week (not at all, a little, quite a bit, or very much). A repeated report going from a little to quite a bit in the course of a week gives relevant information (namely, an increase in pain over time for this patient) but it is challenging to generalize these subjective assessments over patients: what is 'quite a bit' for some patient may feel as 'very much' for one other, depending on their 'baseline' [13]. This makes it difficult to use this information to learn and encode *averages* over individual QoL scores in Bayesian networks *independently* of this baseline and then use these averages to make individual predictions.

Furthermore, when using such information, we are typically interested in change, or even rate of change, rather than absolute value; e.g., in decision making we might search for a possible treatment that slows down a decreasing QoL without compromising disease-specific survival rate [11]. We want to make complex inferences, taking into account time-dependent changes in the relative difference between some baseline (no treatment) and a potential course of action, to advise patients and healthcare professionals on the likely outcome of decisions to be made. This substantially surpasses the well-known algorithmic approaches in Bayesian inferences, like computing posterior distributions or most probable explanations.

To summarize, in order to systematically include quality-of-live in clinical decision support systems, one needs to reason with subjective, but ordered and temporal information ("I'm feeling better today than yesterday"). This implies novel representation, learning, inferring, and explanation of such information in Bayesian networks. In this short paper I will present some initial ideas with respect to representation of subjective, temporal, and ordered information, how such representations can be learned from heterogeneous and subjective QoL data, and how complex inferences (like 'what is the probability that this treatment increases quality-of-life?') can be computed using such representations.

[1] https://www.eortc.org/app/uploads/sites/2/2018/08/Specimen-QLQ-C30-English.pdf

1.1 Related work

We assume that the reader is in general familiar with Bayesian networks and refer to textbooks like [3] for an introduction. *Dynamic* Bayesian networks [12] enhance Bayesian networks (representing a static joint probability distribution) with a notion of time-related dependencies. They are typically modelled as a prior model and a transition distribution; for inference this is typically unpacked into several time-slices where each slice is a (static) Bayesian network and intra-time-slice arcs depict changes over time. In the literature there have been several attempts to model dynamic information in CDSSs using dynamic Bayesian networks [10,17,19]; however, none of these models include QoL information.

Other CDSSs, such as [20], do integrate QoL information (in the form of average FACT-B[2] scores) in the Bayesian network. However, the way this information is encoded makes it difficult to explicitly reason about (desired or expected) outcomes over time. Furthermore, this approach is limited to QoL scores that are summarized from validated tests batteries such as the FACT-B questionnaire.

Subjective Bayesian networks [1], finally, do not encode subjective (non-absolute) *information* but are designed to encode subjective *probability assessments*, using imprecise probabilities. A similar approach to explicitly modelling ignorance about exact probabilities (using credal networks, [2]) has found some application in decision support systems, but also here the focus is on representing sets of probabilities rather than on representing subjective variables.

1.2 Remainder of this paper

The remainder of this paper is organized as follows. We introduce our proposed representation of temporal QoL information in Bayesian networks in Section 2, and illustrate this approach with examples problems in learning (Section 3) and inference (Section 4). We conclude in Section 5.

2 Representation

We make the assumption that the Bayesian networks in question combine patient-specific characteristics (generic aspects, like age and sex, as well as findings, measurements, and observations, like tumour size and type); decision variables, like treatment options, and outcomes, like disease-specific survival rate. A general depiction of the typical relationships between these variables (PC for patient characteristics, T for treatment, and SR for survival rate) is given in Figure 1(a). Note that this is an overly simplified model and in reality (such as the ENDORISK model) there will be more complex interactions, hidden variables, etc. It suffices to make our point, though.

Typically one makes a decision (establishes a value for T) to optimize the outcome SR, taken the patient characteristics into account. Computationally, this

[2] https://www.facit.org/measures/fact-b

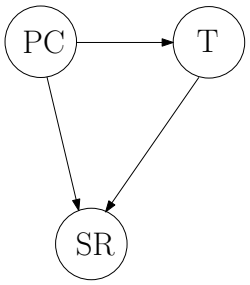

(a) A Bayesian network relating patient's characteristics (PC) with possible treatment (T) and survival rate (SR). This hugely simplified toy example illustrates the key inference to be made using such networks: Given the patient's characteristics, find the treatment that optimizes the expected survival rate.

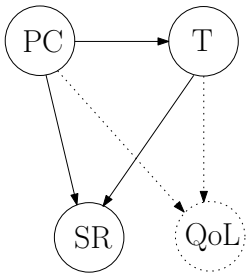

(b) In this adapted network we add QoL considerations in the form of an additional node QoL. Note that we now have a trade-off between optimizing for SR and for QoL, allowing patients and healthcare practitioners to discuss preferences over different outcomes.

Fig. 1: Integrating QoL information in Bayesian decision support systems - a naive approach.

boils down to the problem "Find a treatment T that maximizes $\mathbb{E}(\text{SR} \mid \text{T}, \text{PC})$" where $\mathbb{E}(.)$ indicates the expected value of a variable, in this case, the expected survival rate.

Now, a naive extension to QoL as second outcome variable is depicted in Figure 1(b). Here, QoL is a similar variable as SR, and if this variable can be summarized in some standard scale (like the FACT-B score) this may well serve for some applications. This approach, however, has several drawbacks:

1. QoL here is a static variable; it does not take change (as a result of treatment) into account. Ideally we would like to consider QoL pre-treatment as well as post-treatment, and in many cases also some time after treatment (follow-up), to compare the effects of treatment over time with a baseline with no treatment.
2. The value of QoL is treated as an absolute value, rather than a patient-specific, time-dependent relative value. While these absolute values may be useful for comparisons over cohorts of patients (averaging out individual differences), they are of limited value to support individual decision making.
3. Numerical values for QoL are not always available or useful. They do not facilitate advanced queries such as "find a treatment that slows down decrease of QoL" nor are easily translatable from statements like "I feel way better after treatment than before".

We therefore introduce a novel proposal with three key ingredients that address these drawbacks:

1. We introduce multiple variables (three, in this presentation, but this can be adjusted where needed) QoL_{pre}, QoL_{post}, and $\text{QoL}_{\text{after}}$. All three are dependent on the patient's characteristics, post and after treatment also on treatment; post and after treatment are dependent on the previous state in time. With treatment set to 'no treatment', these variables encode endogenous change in QoL, otherwise they encode the relative effect of the treatment.
2. QoL_{pre} encodes a baseline state pre-treatment; QoL_{post} and $\text{QoL}_{\text{after}}$ encode changes relative to the baseline. In our approach these variables take values from a Likert scale (for example, 'bad', 'OK', and 'fine' for QoL_{pre} and 'worse', 'similar', and 'better' for QoL_{post} and $\text{QoL}_{\text{after}}$). More fine-grained granularity or more specific values can be used where desired.
3. We introduce a *virtual* binary node f_{QoL} which represents a deterministic statement given its parents; $\Pr(f_{\text{QoL}} = T \mid \pi) = 1$ for a specific set of values π of its parents if and only if these values are consistent with the statement. An example of such a statement would be 'Quality of life does not decrease post-treatment'; the corresponding CPT for $\Pr(f_{\text{QoL}} = T \mid \pi)$ evaluates to 1 if and only if the values for QoL_{post} and $\text{QoL}_{\text{after}}$ are 'similar' or 'better'. This virtual node is used for both learning and inference.

Figure 2 captures this idea. Note that f_{QoL} is a virtual node: it is temporally added to the network upon learning or inference as desired, as depicted in Sections 3 and 4.

2.1 Running example

Experienced pain is a notoriously subjective self-reported outcome. To some extent it may be validated, for example by cross-validation with other reported measures [6] or by the physical validation of reported pain with behaviour like moaning, sweating, and asking for medication [5]. However, "*[a]mong patients with the same condition, pain ratings typically cover the entire scale from "no pain" to "the worst pain imaginable."* " [13, p. 231]. In this running example we take inspiration from a study of phantom pain after limb amputation [14]; participants were questioned about pain one day prior, one week after, and six months after scheduled lower limb amputation. In the above study, pain was reported using the Visual Analog Scale (VAS, where participants put a mark on a 0-10 scale with 0 meaning 'no pain at all' and 10 meaning 'pain as bad as it could possibly be'); to illustrate our approach we assume that our data is way more heterogeneous in nature and may range from (three-valued) VAS scores to statements verbally relating pain post-operation to pre-operation.

A partial Bayesian network is shown in Figure 3. The distributions $\Pr(\text{pre})$, $\Pr(\text{post} \mid \text{amp} = T)$ and $\Pr(\text{after} \mid \text{amp} = T)$ are not yet shown as we will learn them from data. The example shows that, when not treated, the expected pain

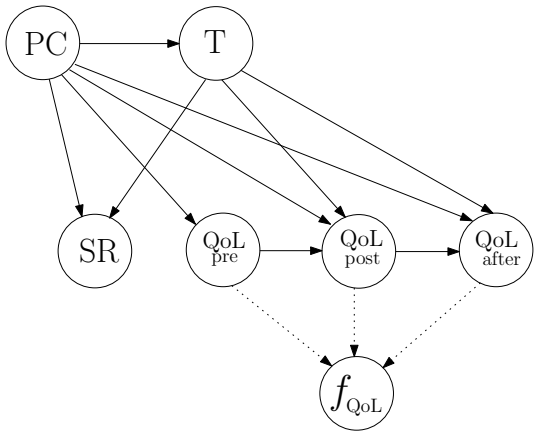

Fig. 2: Essence of the proposed representation. Rather than a single node, QoL is represented as a series of nodes (pre-treatment, post-treatment, and follow-up after treatment). The virtual node f_{QoL} is a deterministic function, representing the 'possible worlds' congruent with a statement or desired outcome; this node is set to true as evidence for learning and inference.

will be more severe. For ease of exposure and training we further simplify the model and assume no arc from QoL_{pre} to QoL_{post} or from QoL_{post} to $\text{QoL}_{\text{after}}$.

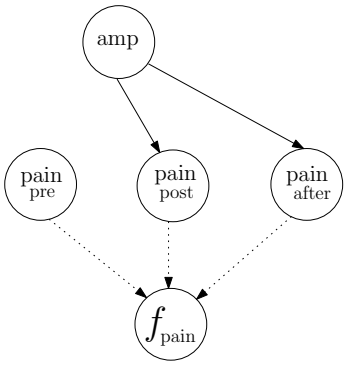

$$
\begin{aligned}
&\Pr(\text{amp} = T) &&= 0.50\\
&\Pr(\text{post} = \text{worse} \mid \text{amp} = F) &&= 0.75\\
&\Pr(\text{post} = \text{similar} \mid \text{amp} = F) &&= 0.20\\
&\Pr(\text{post} = \text{better} \mid \text{amp} = F) &&= 0.05\\
&\Pr(\text{after} = \text{worse} \mid \text{amp} = F) &&= 0.80\\
&\Pr(\text{after} = \text{similar} \mid \text{amp} = F) &&= 0.15\\
&\Pr(\text{after} = \text{better} \mid \text{amp} = F) &&= 0.05
\end{aligned}
$$

Fig. 3: Example network with partial CPT. Note that f_{pain} is a virtual node; we will learn Pr(pre), $\Pr(\text{post} \mid \text{amp} = T)$ and $\Pr(\text{after} \mid \text{amp} = T)$ from data.

3 Learning

As indicated above we want to be able to learn QoL_{pre}, QoL_{post}, and $\text{QoL}_{\text{after}}$ from heterogeneous data, both from actual values (like numerical VAS scores or Likert scale scores) as well as from descriptive verbal statements. For actual values (that map the state space of the QoL variables, if needed after pre-processing) we translate absolute values into relative values and use a vanilla Expectation-Maximization technique. For descriptive values we use the virtual evidence node f_{pain} and treat the QoL variables as hidden variables. For illustration some hypothetical data points are given in Table 1. Note that both the post-treatment and follow-up change are relative to pre-treatment.

	pre-treatment	post-treatment (absolute)	post-treatment (relative)	follow-up (absolute)	follow-up (relative)
p1	severe	moderate	better	moderate	better
p2	severe	severe	similar	moderate	better
p3	moderate	moderate	similar	moderate	similar
p4	severe	severe	similar	moderate	better
p5	moderate	severe	worse	absent	better
p6	'I felt worse after treatment than before, but much better after a while'				
p7	'I was in very much pain pre-treatment, but it improved in time'				
p8	'My pain has not changed, but it has only been a week after treatment'				

Table 1: Small example data set of patients self-reporting pain using a questionnaire before, post, and following up on treatment, as well as three verbal reports of pain dynamics. Absolute reported values are adjusted into a relative value for post and follow up treatment.

For p6, p7, and p8, we translate the statements to a deterministic CPT for the virtual evidence variable f_{pain}. We set $\Pr(f_{\text{pain}} = T \mid \text{pre}, \text{post}, \text{after}) = 1$ if and only if the statement is consistent with the values for pre, post, and after. Note there is sometimes some (subjective) interpretation necessary. For example, for p7 we assumed pain pre-treatment was severe (not moderate), we are indifferent about the change briefly after treatment, and assumed improvement at follow-up.

We use the Expectation-Maximization algorithm to augment the data points for p6, p7, and p8 (treating pre, post, and after as hidden variables with a different CPT for f_{pain} for each patient). The weighted data points are then normalized and the resulting probabilities are listed in Figure 4.

4 Inference

In the previous section we saw how virtual evidence nodes can help translate verbal statements into consistent values for pre, post, and after, facilitating learning

p5	post = better			post = similar			post = worse		
	after = b	after = s	after = w	after = b	after = s	after = w	after = b	after = s	after = w
pre = bad	0	0	0	0	0	0	1	0	0
pre = OK	0	0	0	0	0	0	1	0	0
pre = fine	0	0	0	0	0	0	1	0	0

p6	post = better			post = similar			post = worse		
	after = b	after = s	after = w	after = b	after = s	after = w	after = b	after = s	after = w
pre = s	0	0	0	0	0	0	0	0	1
pre = m	0	0	0	0	0	0	0	0	1
pre = a	0	0	0	0	0	0	0	0	1

p7	post = better			post = similar			post = worse		
	after = b	after = s	after = w	after = b	after = s	after = w	after = b	after = s	after = w
pre = s	1	0	0	1	0	0	1	0	0
pre = m	0	0	0	0	0	0	0	0	0
pre = a	0	0	0	0	0	0	0	0	0

p8	post = better			post = similar			post = worse		
	after = b	after = s	after = w	after = b	after = s	after = w	after = b	after = s	after = w
pre = s	0	0	0	1	1	1	0	0	0
pre = m	0	0	0	1	1	1	0	0	0
pre = a	0	0	0	1	1	1	0	0	0

Table 2: Deterministic CPT for f_{pain} for p6, p7, and p8.

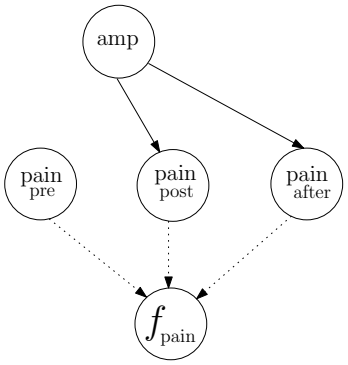

$$
\begin{aligned}
&\Pr(\text{amp} = T) && = 0.50\\
&\Pr(\text{pre} = \text{severe}) && = 0.58\\
&\Pr(\text{pre} = \text{moderate}) && = 0.34\\
&\Pr(\text{pre} = \text{absent}) && = 0.08\\
&\Pr(\text{post} = \text{worse} \mid \text{amp} = F) && = 0.75\\
&\Pr(\text{post} = \text{similar} \mid \text{amp} = F) && = 0.20\\
&\Pr(\text{post} = \text{better} \mid \text{amp} = F) && = 0.05\\
&\Pr(\text{post} = \text{worse} \mid \text{amp} = T) && = 0.08\\
&\Pr(\text{post} = \text{similar} \mid \text{amp} = T) && = 0.33\\
&\Pr(\text{post} = \text{better} \mid \text{amp} = T) && = 0.59\\
&\Pr(\text{after} = \text{worse} \mid \text{amp} = F) && = 0.80\\
&\Pr(\text{after} = \text{similar} \mid \text{amp} = F) && = 0.15\\
&\Pr(\text{after} = \text{better} \mid \text{amp} = F) && = 0.05\\
&\Pr(\text{after} = \text{worse} \mid \text{amp} = T) && = 0.04\\
&\Pr(\text{after} = \text{similar} \mid \text{amp} = T) && = 0.17\\
&\Pr(\text{after} = \text{better} \mid \text{amp} = T) && = 0.79
\end{aligned}
$$

Fig. 4: Example network after learning.

a model from heterogeneous data. Interestingly, a similar approach can be used for facilitating complex inference queries. For example, we may want to compute

the probability that a given treatment slows down or reverts a decrease of QoL post treatment. In our simple example, this is the case if the value for after is better than or equal to the value for post. The corresponding CPT is given in Table 3.

p6	post = b			post = s			post = w		
	after = b	after = s	after = w	after = b	after = s	after = w	after = b	after = s	after = w
pre = s	1	0	0	1	1	0	1	1	0
pre = m	1	0	0	1	1	0	1	1	0
pre = a	1	0	0	1	1	0	1	1	0

Table 3: Deterministic CPT for f_{pain} consistent with the statement 'Treatment slows down or reverts a decrease of QoL.

We can then compute, for example, $\Pr(f_{\text{pain}} = T \mid \text{pre} = \text{moderate}, \text{amp} = T)$ in the trained model above and find that $\Pr(f_{\text{pain}} = T \mid \text{pre} = \text{moderate}, \text{amp} = T) = 0.86$.

In more complicated models such functions can facilitate elaborate queries. For example, assume that multiple treatments exist and that QoL_{pre}, QoL_{post}, and $\text{QoL}_{\text{after}}$ all have a five point Likert scale (ranging from $--$ to $++$). Using this approach we can then compute complex queries like 'Find a treatment that slows down the decrease in QoL.' or 'Find a treatment that has the best expected QoL while having an expected disease-specific survival rate of at least 3 years.' Of course, this all depends on the availability of a correct mapping from the variables of interest to f_{pain}.

5 Conclusion

In this short paper we presented some early ideas on how to represent, learn, and reason with relative and temporal information in Bayesian networks, allowing for complex inferences about desired or expected changes in quality-of-life of patients. Crucial aspects of this approach are QoL-variables before and after treatment where the latter encode a *change* relative to the pre-treatment baseline. Complex dynamics can be encoded with multiple time slices; this allows for inferences about change or even rate of change due to some intervention, such as the application of a treatment.

Important in this approach is the introduction of a virtual evidence variable while learning or during inference; this virtual evidence variable is a deterministic variable that encodes the QoL values consistent with a particular *statement*, i.e., a mapping $f(\text{statement}) \rightarrow$ possible worlds.

Different variations of this proposed formalism are possible. In our running example we used three values per QoL variable; other approaches may use a Likert scale of four or five variables. We would advise to be careful with interpreting numerical values such as a VAS scale unless patient scores are validated and normalized, as the average score may suggest an unwarranted precision. We used a

concrete baseline (i.e., severe, moderate, or absent pain); one can also decide to model change relative to an unspecified baseline ('better than before'). In our running example we left out endogenous change (i.e., no arcs between the QoL variables); one might of course go in the other direction and assume interaction effects between pre-treatment and post-treatment at follow-up measurement. Finally we encoded change relative to pre-treatment for both post-treatment and follow-up; one might also encode change relative to the previous measurement moment (follow-up vs. post-treatment).

Obvious future work is to explore actual QoL data sets to study the training of concrete networks such as ENDORISK; furthermore it would be of interest to study the different variations for baseline values etc. as indicated above. More fundamental possible work seeks to optimize inference given the deterministic structure of f_{QoL} cfm. [4], to decompose complex inference queries to simpler logical structures (using multiple virtual evidence nodes) [8], and to explore explanation of QoL results.

Acknowledgments. Several of the ideas put forward in this paper were discussed at lab meetings of the PGM group at Radboud University and I am indebted for useful feedback offered by group members, particularly Wim Wiegerinck. I also thank the reviewers for there insightful comments and suggestions.

Disclosure of Interests. The author has no competing interests to declare that are relevant to the content of this article.

References

1. Braines, D., Thomas, A., Kaplan, L., Şensoy, M., Bakdash, J.Z., Ivanovska, M., Preece, A., Cerutti, F.: Subjective Bayesian networks and human-in-the-loop situational understanding. In: Graph Structures for Knowledge Representation and Reasoning: 5th International Workshop, GKR 2017, Melbourne, VIC, Australia, August 21, 2017, Revised Selected Papers 5. pp. 29–53. Springer (2018)
2. Cozman, F.G.: Credal networks. Artificial intelligence **120**(2), 199–233 (2000)
3. Darwiche, A.: Modeling and reasoning with Bayesian networks. Cambridge university press (2009)
4. Evers, S., Lucas, P.J.F.: Marginalization without summation exploiting determinism in factor algebra. In: Liu, W. (ed.) Proceedings of ECSQARU 2011. LNCS, vol. 6717, pp. 251–262. Springer (2011). https://doi.org/10.1007/978-3-642-22152-1_22
5. Fletcher, G.S.: Clinical epidemiology: the essentials. Lippincott Williams & Wilkins (2019)
6. Hidalgo, B., Goodman, M.: Validation of self-reported measures in health disparities research. Journal of biometrics & biostatistics **3**(7) (2012)
7. Kasztura, M., Richard, A., Bempong, N.E., Loncar, D., Flahault, A.: Cost-effectiveness of precision medicine: a scoping review. International journal of public health **64**, 1261–1271 (2019)

8. Kwisthout, J.: Most inforbable explanations: Finding explanations in Bayesian networks that are both probable and informative. In: van der Gaag, L. (ed.) Proceedings of the Twelfth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty. LNAI, vol. 7958, pp. 328–339. Springer-Verlag (2013)
9. Kyrimi, E., McLachlan, S., Dube, K., Neves, M.R., Fahmi, A., Fenton, N.: A comprehensive scoping review of Bayesian networks in healthcare: Past, present and future. Artificial Intelligence in Medicine **117**, 102108 (2021)
10. Ltifi, H., Trabelsi, G., Ayed, M.B., Alimi, A.M.: Dynamic decision support system based on Bayesian networks. International Journal of Advanced Research in Artificial Intelligence **1**(1) (2012). https://doi.org/10.14569/IJARAI.2012.010105, http://dx.doi.org/10.14569/IJARAI.2012.010105
11. Movsas, B., Vikram, B., Hauer-Jensen, M., Moulder, J.E., Basch, E., Brown, S.L., Kachnic, L.A., Dicker, A.P., Coleman, C.N., Okunieff, P.: Decreasing the adverse effects of cancer therapy: National Cancer Institute guidance for the clinical development of radiation injury mitigators. Clinical Cancer Research **17**(2), 222–228 (2011)
12. Murphy, K.P.: Dynamic Bayesian networks: representation, inference and learning. University of California, Berkeley (2002)
13. Nielsen, C.S., Staud, R., Price, D.D.: Individual differences in pain sensitivity: measurement, causation, and consequences. The journal of pain **10**(3), 231–237 (2009)
14. Nikolajsen, L., Ilkjær, S., Jensen, T.S.: Relationship between mechanical sensitivity and postamputation pain: a prospective study. European journal of pain **4**(4), 327–334 (2000)
15. Reijnen, C., Gogou, E., Visser, N.C., Engerud, H., Ramjith, J., Van Der Putten, L.J., Van de Vijver, K., Santacana, M., Bronsert, P., Bulten, J., et al.: Preoperative risk stratification in endometrial cancer (ENDORISK) by a Bayesian network model: A development and validation study. PLoS medicine **17**(5), e1003111 (2020)
16. Revicki, D.A., Osoba, D., Fairclough, D., Barofsky, I., Berzon, R., Leidy, N., Rothman, M.: Recommendations on health-related quality of life research to support labeling and promotional claims in the United States. Quality of Life Research **9**, 887–900 (2000)
17. Rose, C., Smaili, C., Charpillet, F.: A dynamic Bayesian network for handling uncertainty in a decision support system adapted to the monitoring of patients treated by hemodialysis. In: 17th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'05). pp. 5–pp. IEEE (2005)
18. Sutton, R.T., Pincock, D., Baumgart, D.C., Sadowski, D.C., Fedorak, R.N., Kroeker, K.I.: An overview of clinical decision support systems: benefits, risks, and strategies for success. NPJ digital medicine **3**(1), 17 (2020)
19. Van Gerven, M., Taal, B., Lucas, P.: Dynamic Bayesian networks as prognostic models for clinical patient management. Journal of biomedical informatics **41**(4), 515–529 (2008)
20. Xu, S., Thompson, W., Ancoli-Israel, S., Liu, L., Palmer, B., Natarajan, L.: Cognition, quality-of-life, and symptom clusters in breast cancer: using Bayesian networks to elucidate complex relationships. Psycho-oncology **27**(3), 802–809 (2018)
21. Zafar, S.Y., Alexander, S.C., Weinfurt, K.P., Schulman, K.A., Abernethy, A.P.: Decision making and quality of life in the treatment of cancer: a review. Supportive Care in Cancer **17**, 117–127 (2009)