Bachelor thesis

Artificial Intelligence

Radboud University

Comparing quantified MAP-dependence to other measures of relevance in Bayesian networks

Author:
Merlijn van Elteren
s1021267

Supervisor:
dr. J.H.P Kwisthout
Artificial Intelligence
j.kwisthout@donders.ru.nl

Second reader:
prof. dr. I.J.E.I. van Rooij
Artificial Intelligence
i.vanrooij@donders.ru.nl

June 20, 2022

Abstract

In the fields of explainable AI and Bayesian networks, a recently introduced concept is a measure called *MAP-independence*. It is meant to assist in the justification of decisions made using decision support systems, by identifying irrelevant intermediate variables in a decision problem. In this paper, we argue that the binary nature of this measure is too crude, and may lead to a large set of relevant variables, some of which only change the *Maximum A Posteriori* outcome for very specific and maybe unlikely observations. We hypothesise that *"quantifying"* the measure (meaning *"to change it from a binary measure to a measure on the [0,1] scale"*) would allow us to identify differences in how relevant each variable in a set of relevant variables is. We name this new measure *quantified MAP-dependence*. We make an implementation which is, to the best of our knowledge, the first implementation of both MAP-independence and quantified MAP-dependence. Furthermore, we apply these measures to the ALARM network. For comparison, we also apply an older measure of relevance named *intrinsic relevance*. Based on the results, we conclude the following: Firstly, that our hypothesis about the usefulness of quantifying MAP-independence is true. Secondly, that compared to intrinsic relevance, quantified MAP-dependence describes a fundamentally different interpretation of what makes an intermediate variable relevant.

Contents

Chapter 1

Introduction

Imagine you are in the hospital, presenting life-threatening symptoms. However, you have no idea what might be causing them. The doctor runs tests, attempting to identify the cause. Yet in the end, they are unsure which out of a few possible illnesses is most likely afflicting you, and the different possibilities require different treatments...

Though somewhat extreme, this is a good example of a situation in which an effective *decision support system (DSS)* would be very useful. Underlying such a system, we may find a statistical model of the specific situation, which the DSS uses to infer the best explanation of the available evidence. Specifically, a *Bayesian network (BN)*, as defined in [Chapter 2], is a type of statistical model which may be applied to different public domains. For the example we described, we could use a BN as a medical diagnostic model [1]. BNs are also applied in other situations, for example as weather prediction models [2], or even to predict the optimal verdict given evidence in court [3].

Returning to our example, we can imagine, both from the perspective of the doctor and that of the patient in question, that simply entering some data and receiving the "most likely" diagnosis may not be enough for us to trust that this diagnosis is accurate. Even assuming we expand this output to a list of the most likely illnesses with their respective likelihoods, we might wonder how the model came to these conclusions. It has been shown that in order for the user to trust and effectively utilise a DSS in situations like this, where important decisions are aided by BNs or other forms of machine learning, the system must be able to explain its decisions [4].

The 2016 changes to the European Union's General Data Protection Regulations reflected this need for being able to explain decisions made using automated systems, by introducing the concept commonly referred to as the "right to an explanation" [5]. This has given rise to the field of *eXplainable Artificial Intelligence (XAI)*. In [6], researchers investigated how an explanation should ideally be structured, based on information from the research areas of pedagogy, story-telling, argumentation, programming, trust-building and gamification. They applied this structure to XAI explanations, explicating the building blocks which an XAI explanation should ideally contain. An AI system which could offer explanations at the level described in that paper, in terms of the "notions of explainability" for XAI [7], would be a truly "explainable system". For context, the notions of explainability are specified in the table below.

Rank	Notion	Meaning
1	opaque system	A system offering no insights into its algorithmic mechanisms; A "black box".
2	interpretable system	A system where users can mathematically analyse its algorithmic mechanisms.
3	comprehensible system	A system which emits symbols enabling user-driven explanations of how a conclusion is reached.
4	explainable system	A system where automated reasoning is central to output crafted explanations without requiring human post processing as final step of the generative process.

Figure 1: The notions of explainability as defined in [7].

However, for BNs, which fit into the "interpretable systems", recent research only aims at turning it into a "comprehensible system". Similarly, most of the latest XAI research in general focuses on turning an "opaque system" into a "comprehensible system", thus explaining what happens in the "black box". In a recent analysis [8] of the effects of some state-of-the-art XAI techniques for generating explanations, the authors defined three "desiderata" of AI explanations; understanding, uncertainty awareness, and trust. However, they showed that none of the analysed techniques conclusively helped with all the criteria. In fact, most of them were found not to help with any of the criteria at all, unless the user felt that they already had some expertise in the relevant AI concepts. As such, there is still work to be done.

Coming back to Bayesian networks, a recently introduced concept is *"MAP-independence"* [9], as defined in [Chapter 2]. When using *Maximum A Posteriori* (MAP, [Chapter 2]) calculations over a Bayesian network, this new measure is used to identify which variables of the network were (ir)relevant to the calculation. We can use this information to explain and justify decisions based on such calculations, effectively turning our originally "interpretable system" into a "comprehensible system".

However, we argue that the all-or-nothing nature of this measure of (ir)relevance may be an obstacle to providing the best possible explanation in every circumstance. MAP-independence has a clear underlying interpretation of what makes a variable relevant, but the measure can only be used to decide whether a variable is completely irrelevant (MAP-independent) or not (MAP-dependent). A quantified version of the measure, as defined in [Chapter 3, Section 1], could be used to measure relevance more sharply. Assume we are measuring on the [0,1] scale. Clearly, if one variable is measured to have a relevance of 0.2 while an other variable is measured to have a relevance of 0.8, this implies quite a large difference in relevance between the variables. However, this difference is not captured by MAP-(in)dependence. Even if we decide to use all the MAP-dependent variables in an explanation, quantifying the measure would allow us to explicate the differences in relevance between variables. Furthermore, we imagine some users might find it easier to form a clear and understandable explanation using, for some number n, the n most relevant variables. MAP-(in)dependence would not allow them to do so, instead restricting them to either using all variables which are relevant, or making an arbitrary decision on which n of those variables they should use. In contrast, using a quantified version of MAP-dependence, this would be easy to accomplish. As such, quantified MAP-dependence may allow us to adapt the level of detail in our explanations to different people receiving the explanation.

To the best of our knowledge, as MAP-independence is such a recent concept, it has not yet been implemented as an algorithm or practically applied to anything except some small example networks, as in [9]. Therefore, we make the first implementation [Chapter 3, Section 2] of both the original MAP-independence, as well as our proposed quantified version. Furthermore, we will apply these measures to MAP-calculations over a well-known benchmark Bayesian network called the *ALARM network* [10], following the methods specified in [Chapter 4]. We will also implement and apply an older measure of relevance called *"intrinsic relevance"* [11], also defined in [Chapter 2]. Since intrinsic relevance can be used for the same purpose as MAP-independence and its quantified version, we will be comparing the results [Chapter 5] provided by applying these different measures to the ALARM network. Overall conclusions are drawn in [Chapter 6].

In short, the main contributions of this paper are as follows:

- We expand upon Kwisthout's idea of MAP-independence, originally a binary (yes or no) concept, by turning it into a quantifiable measure, measured on the [0,1] scale.

- We implement the measure for the first time, and compare the results it produces on a well-known benchmark network, to those of a previously established measure of relevance.

Chapter 2

Preliminaries

In this chapter we provide some preliminary information, and introduce the notational conventions we use. We find both to be essential to understanding further content of this paper. For more background information on Bayesian networks, the reader can consider textbooks like [12].

In the preliminaries section of the main article [9] which we build upon, Kwisthout introduces some key concepts and relevant notational conventions, which we also use:

> A Bayesian network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Pr)$ is a probabilistic graphical model that succinctly represents a joint probability distribution $\Pr(\mathbf{V}) = \prod_{i=1}^{n} \Pr(V_i \mid \pi(V_i))$ over a set of discrete random variables $\mathbf{V}$. $\mathcal{B}$ is defined by a directed acyclic graph $\mathbf{G}_{\mathcal{B}} = (\mathbf{V}, \mathbf{A})$, where $\mathbf{V}$ represents the stochastic variables and $\mathbf{A}$ models the conditional (in)dependencies between then, and a set of parameter probabilities Pr in the form of conditional probability tables (CPTs). In our notation $\pi(V_i)$ denotes the set of parents of a node V_i in $\mathbf{G}_{\mathcal{B}}$. We use upper case to indicate variables, lower case to indicate a specific value of a variable, and boldface to indicate sets of variables respectively joint value assignments to such a set. $\Omega(V_i)$ denotes the set of value assignments to V_i, with $\Omega(\mathbf{V_a})$ denoting the set of joint value assignments to the set $\mathbf{V_a}$. One of the key computational problems in Bayesian networks is the problem to find the most probable explanation for a set of observations, i.e., a joint value assignment to a designated set of variables (the explanation set) that has maximum posterior probability given the observed variables (the joint value assignment to the evidence set) in the network. If the network includes variables that are neither observed nor to be explained (referred to as intermediate variables) this problem is typically referred to as MAP. We use the following formal definition:

MAP
Instance: A Bayesian network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Pr)$, where $\mathbf{V}(\mathbf{G}_{\mathcal{B}})$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$, a set of intermediate nodes $\mathbf{I}$, and an explanation set $\mathbf{H}$.
Output: A joint value assignment $\mathbf{h}^*$ to $\mathbf{H}$ such that for all joint value assignments $\mathbf{h}'$ to $\mathbf{H}$, $\Pr(\mathbf{h}^* \mid \mathbf{e} \geq \Pr(\mathbf{h}' \mid \mathbf{e})$.

As quoted from [9, p.5].

After introducing these concepts, Kwisthout goes on to introduce his idea, MAP-independence; What it is, why it is relevant, and how it differs from related concepts. For a complete understanding of the concept, we wholeheartedly recommend reading section 3 [9, pp. 6-9], titled MAP-independence. However, in order to keep this chapter to the point, we continue at the start of the Kwisthout's next section, wherein he formalises MAP-independence:

The computational problem of interest is to decide upon the set $\mathbf{I}^+$, the relevant variables that contribute to establishing the best explanation $\mathbf{h}^*$ given the evidence $\mathbf{e}$. In order to establish $\mathbf{I}^+$ we need to decide the following sub-problem: given $\mathbf{R} \subseteq \mathbf{I}$: is $\mathbf{h}^*$ MAP-independent from $\mathbf{R}$ given $\mathbf{e}$? We formalise this problem as below.

MAP-INDEPENDENCE
Instance: A Bayesian network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Pr)$, where $\mathbf{V}(\mathbf{G})$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$, a non-empty explanation set $\mathbf{H}$ with a joint value assignment $\mathbf{h}^*$, a non-empty set of nodes $\mathbf{R}$ for which we want to decide MAP-independence relative to H, and a set of intermediate nodes $\mathbf{I}$.
Question: Is $\forall_{\mathbf{r}\in\Omega(\mathbf{R})}\text{argmax}_{\mathbf{H}}\Pr(\mathbf{H}, \mathbf{R} = \mathbf{r} \mid \mathbf{e}) = \mathbf{h}^*$?

Observe that the complement problem MAP-DEPENDENCE is defined similarly with *yes*- and *no*-answers reversed.

In this formal definition we allow for interaction in the set R. The related problem Weak MAP-independence, where we consider MAP-independence for each singleton variable on its own, is defined as follows:

WEAK MAP-INDEPENDENCE
Instance: As in MAP-INDEPENDENCE .

Question: Is $\forall_{R \in \mathbf{R}} \forall_{r \in \Omega(R)} \text{argmax}_{\mathbf{H}} \Pr(\mathbf{H}, R = r | \mathbf{e}) = \mathbf{h}^*$?

As quoted from [9, pp. 9-10].

Note that Kwisthout also formally defined a problem closely related to MAP-independence, which he called *weak MAP-independence*. It is important to know that we have restricted our research to exclusively consider instances where the set $\mathbf{R}$ has size one. This essentially means that we only calculate the relevance measures over different singleton variables R. We find it astute to mention that, due to this restriction, MAP-independence and weak MAP-independence are effectively equal within the context of our paper. From this point on, we will only explicitly differentiate between weak MAP-independence and MAP-independence, if we are specifically comparing the two concepts. Otherwise, we will exclusively use the term *MAP-independence*, because that concept is central to our research.

The final concept we define in this chapter, is that of *intrinsic relevance*. It was formalised by Kwisthout, based on Druzdel and Suermondt's [13] definition of relevance of variables in a Bayesian network. For a deeper understanding, we recommend reading section 3.1, titled Relevance, from [11, pp. 3-4]. The essential information, however, is condensed in the formal definition of intrinsic relevance, as provided below:

INTRINSIC RELEVANCE [11]
Let $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Pr)$ be a Bayesian network partitioned into evidence nodes $\mathbf{E}$ with joint value assignment $\mathbf{e}$, intermediate nodes $\mathbf{I}$, and an explanation set $\mathbf{H}$. Let $R \in \mathbf{I}$, and let $\Omega(\mathbf{I} \setminus \{R\})$ denote the set of joint value assignments to the intermediate variables other than R. The *intrinsic relevance* of R is the fraction of joint value assignments $\mathbf{r}$ in $\Omega(\mathbf{I} \setminus \{R\})$ for which $\text{argmax}_h \Pr(\mathbf{h}, \mathbf{e}, \mathbf{i}, r)$ is not identical for all $r \in \Omega(R)$.

Note that this measure of relevance is only defined over a singleton variable, which we have called R in this paper in order to improve readability in relation to the definition of MAP-independence, compared to the original name I as in [11]. The fact that this measure is only defined over a singleton variable motivates the previously explicated restriction to our research, which is to only consider sets $\mathbf{R}$ of size one. Without this restriction, we could not compare the measures of MAP-independence and intrinsic relevance.

Chapter 3

Quantified MAP-dependence

In [Chapter 1], we exemplified the kind of situation in which MAP-independence might be used, and argued how we might benefit from quantifying this binary measure. However, we have not yet stated exactly how we would quantify the measure. This is explained in the first section of this chapter. In the second section, we highlight essential parts of our implementation of the relevance measures.

3.1 Definition

As we can see in [Chapter 2], MAP-independence is defined as a decision problem with an instance and a decision question. The decision question relies upon a for-all-statement, essentially asking whether the outcome $\mathbf{h}^*$ of MAP would be the same regardless of which value assignment $\mathbf{r}$ is given to a sub-set $\mathbf{R}$ of the set $\mathbf{I}$ of intermediate variables. If this is the case, then $\mathbf{R}$ is MAP-independent, and we consider the variables in $\mathbf{R}$ to be irrelevant. But if even a single value assignment of $\mathbf{R}$ does cause a change in the outcome of MAP, then $\mathbf{R}$ is MAP-dependent, and we consider the variables in $\mathbf{R}$ to be relevant.

We observe a clear underlying interpretation of what makes a variable relevant. We think that this interpretation can also be used to define a quantified measure for exactly *how relevant* a relevant variable is. Because if the difference between relevant and irrelevant is captured by asking the question "Does assigning any of the possible values change the outcome of MAP?", then we think relevance could be quantified by asking the question "For what fraction of the possible values does the outcome of MAP change?". Based on this interpretation, we formally define quantified MAP-dependence as follows:

QUANTIFIED MAP-DEPENDENCE
Let $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Pr)$ be a Bayesian network, where $\mathbf{V}(\mathbf{G})$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$, a set of intermediate nodes $\mathbf{I}$ and a non-empty explanation set $\mathbf{H}$ with a joint value assignment $\mathbf{h}^*$. Let $\mathbf{R}$ be a non-empty subset of $\mathbf{I}$, and let $\Omega(\mathbf{R}\}$ denote the set of joint value assignments to the intermediate variables $\mathbf{R}$. The *quantified MAP-dependence* of $\mathbf{R}$ relative to $\mathbf{H}$, denoted as $\mu(\mathbf{R})$, is the fraction of joint value assignments $\mathbf{r}$ in $\Omega(\mathbf{R}\}$ for which $\text{argmax}_{\mathbf{H}}\Pr(\mathbf{H}, \mathbf{R} = \mathbf{r} \mid \mathbf{e})$ is not equal to $\mathbf{h}^*$.

Observe that a set $\mathbf{R}$ with $\mu(\mathbf{R}) = 0$ is MAP-independent, while a set $\mathbf{R}$ with $\mu(\mathbf{R}) > 0$ is MAP-dependent.

Recall from [Chapter 2], that we explicated a restriction which we have put on our research, which is to exclusively consider instances where the set $\mathbf{R}$ has size one. This restriction also applies to quantified MAP-dependence, for the same reason.

3.2 Implementation

Essential parts of the implementation we use to perform calculations (of quantified MAP-dependence, and also of MAP-independence and intrinsic relevance), will be shown in this section. As the implementation is an essential part of our *method* of gathering results (further specified in [Chapter 4]), this is essential to the reproducibility of this research.

For this research, we started out with the Java code for a previously implemented variable elimination algorithm. This implementation already included a method for reading in an existing Bayesian network from a .bif file. Furthermore, it included classes used to represent variables and their (conditional) probability distributions, as well as perform typical variable elimination calculations, such as factor-reduction, -marginalisation and -product. Overall, it was a functional variable elimination algorithm, able to perform MAP calculations over a Bayesian network given some input.

The first set of changes to this preexisting implementation were written in cooperation with Simon Janssen and Luuk Jacobs. As a group, we added the *RelevanceMeasurer* class and wrote the method *solveMapIndependence* as well as some supporting methods. To the best of the group's knowledge, this was the first implementation of MAP-independence. After MAP-independence was successfully implemented, each of us made further changes separately, fitting the needs of our individual projects.

For this project, further additions include the addition of the method *solveIntrinsicRelevance*, project-specific changes to the method *solveMapIndependence*, and a number of quality improvements. Although everything is documented in the complete code, as linked in [Appendix A], we show the two important methods in the form of pseudo-code. This is for the purpose of completeness without having to consult the appendices. See the pseudo-code algorithms below:

Algorithm 1 SOLVEMAPINDEPENDENCE

Input: A Bayesian network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Pr)$, where $\mathbf{V}(\mathbf{G})$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$, a set of intermediate nodes $\mathbf{I}$ and a non-empty explanation set $\mathbf{H}$. A set F of factors representing the network.
Output: Two lists. The first contains the quantified MAP-dependence value for each variable in $\mathbf{I}$, and the second contains a Boolean indicating whether a variable is MAP-independent or not, for each variable in $\mathbf{I}$.

```
listOfQMDs ← new List
listOfMIs ← new List
F.reduce(e)
F_copy ← copy(F)
h* ← solveMAP(F_copy)
for all R ∈ I do
    counter ← 0
    total ←| Ω(R) |
    for all r ∈ Ω(R) do
        F_copy ← copy(F)
        F_copy.reduce(r)
        h′ ← solveMAP(F_copy)
        if h′ ≠ h* then
            counter += 1
        end if
    end for
    QMD = counter / total
    listOfQMDs.add(QMD)
    listOfMIs.add(QMD==0)
end for
return listOfQMDs, listOfMIs
```

Algorithm 2 SOLVEINTRINSICRELEVANCE

Input: A Bayesian network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Pr)$, where $\mathbf{V}(\mathbf{G})$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$, a set of intermediate nodes $\mathbf{I}$ and a non-empty explanation set $\mathbf{H}$. A set F of factors representing the network.

Output: A list containing, for each variable in $\mathbf{I}$, its intrinsic relevance.

```
listOfIRs ← new List
F.reduce(e)
for all R ∈ I do
    assignmentsOfR ← list(Ω(R))
    counter ← 0
    total ←| Ω(I \ {R}) |
    for all i ∈ Ω(I \ {R}) do
        F_copy1 ← copy(F)
        F_copy1.reduce(i)
        F_copy2 ← copy(F_copy1)
        F_copy2.reduce(assignmentsOfR.get(0))
        hFirst = solveMAP(F_copy2)
        for int j = 1; j < assignmentsOfR.size(); j++ do
            F_copy2 ← copy(F_copy1)
            F_copy2.reduce(assignmentsOfR.get(j))
            hOther = solveMAP(F_copy2)
            if hOther ≠ hFirst then
                counter += 1
                break
            end if
        end for
    end for
    IR = counter / total
    listOfIRs.add(IR)
end for
return listOfIRs
```

Chapter 4

Method

In this research, our aim is to compare the different relevance measurements which we have discussed (MAP-independence, quantified MAP-dependence and intrinsic relevance). Specifically, we hypothesise that "quantifying" the measure (meaning "to change it from a binary measure to a measure on the [0,1] scale") would allow us to identify differences in how relevant each variable in a set of relevant variables is. We have no specific hypothesis about how quantified MAP-dependence compares to intrinsic relevance, but simply aim to compare the two. In order to compare the different relevance measurements, we need to find directly comparable results. We will find these results by performing these measurements on a set of instances of a Bayesian network.

The implementation, as discussed in the previous chapter, allows us to perform the relevance measurements on a given instance. Practically speaking, we can *use* these different measurements to classify (sets of) variables as relevant or irrelevant to the calculation of MAP over the given instance. However, it is important to remember that the *end goal* is to explain and justify the decision offered by the decision support system. Although we suspect this purpose can be fulfilled using those variables which we classify as relevant, the classification itself is not the main focus, but rather the explanation we can form using the classification. Therefore, we intend to apply the measurements to meaningful instances, and take a qualitative approach to the analysis of the results, rather than for example working with randomised instances of Bayesian networks, as we would do when taking a more statistical, quantitative approach. In doing so, we hope to preserve the meaningfulness of the network instance, enabling us to make meaningful inferences based on the results we find.

With the previous statement in mind, the Bayesian network we have selected is the well-known ALARM network [10]. It is a network from the medical domain, containing information relating to eight medical diagnoses, which we refer to as the "*hypothesis variables*". Each of the hypothesis variables is linked, either directly or through one or more of the thirteen *intermediate variables*, to one or more of the sixteen measurable findings, which we refer to as "*evidence variables*". As stated in [10, pp. 249], "ALARM contains statistical data on prior probabilities, logical conditional probabilities computed from equations relating variables, and a number of subjective assessments." Because the network was designed by domain experts, and seems to have stood the test of scientific scrutiny, we will assume that it accurately models the situation it describes. That the network is an accurate model is important to us, since otherwise the network would not be as meaningful.

For the next two paragraphs, keep the following in mind: An *input* to our implementation, given a certain network, essentially consists of two things; A set $\mathbf{E}$ of evidence variables with a joint value assignment $\mathbf{e}$, and a non-empty set $\mathbf{H}$ of hypothesis variables. Such an input implicitly defines the set $\mathbf{I}$ of intermediate variables (as all the remaining variables of the network). The non-empty set $\mathbf{R}$ of intermediate variables, for which we calculate the relevance measures, is restricted to always be of size one, as explained in [Chapter 2]. For each input, we select the singleton variable R for an instance by looping over the set $\mathbf{I}$. This essentially means that, for a given input, we create one instance for each variable in $\mathbf{I}$. We calculate our relevance measures over each of these instances.

As we are taking a qualitative approach, it would not make sense to calculate our measurements over all possible inputs. Instead, we will select 20 inputs to analyse. However, we suspect that taking 20 random inputs out of all the possible inputs is not likely to result in the 20 most meaningful inputs. Thus, we have decided to impose some restrictions. As mentioned previously, the ALARM network is defined with clearly determined sets of evidence, intermediate and hypothesis variables. We will not deviate from these predetermined set-allocations, as this would move away from the intended use of the network and therefore likely reduce the meaningfulness of an input. This means that for our research, $\mathbf{H}$ will always consist of the eight medical diagnoses, $\mathbf{I}$ will always consist of the thirteen intermediate variables, and $\mathbf{E}$ will always consist of the sixteen measurable findings, as defined in [10].

However, this leaves one part of the input to be decided; the joint value assignment $\mathbf{e}$ of the set of evidence variables. As stated previously, we want to find the 20 most meaningful inputs. Since $\mathbf{e}$ is the only part of the input that remains variable, we essentially want to find the 20 most meaningful joint value assignments $\mathbf{e}$. We define the most meaningful value assignments as being the 20 value assignments that are most likely to occur in real life. Considering that we assume the network accurately models real life for the described situation, "most likely to occur in real life" is the same as "most likely to occur in the network". In order to find the most likely value assignments according to the network, we use variable elimination to calculate[1] the marginal prior of every possible value assignment. We then sort the list of marginal priors based on their probabilities, and select the top 20. The 20 most likely value assignments and their respective marginal prior probabilities are shown in the following tables:

	HISTORY	CVP	PCWP	HRBP	HREKG	HRSAT	TPR	EXPCO2	probability
1	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	NORMAL	LOW	0.02676
2	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	LOW	LOW	0.02489
3	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	HIGH	LOW	0.02175
4	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	NORMAL	LOW	0,01737
5	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	NORMAL	LOW	0,01687
6	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	LOW	LOW	0,01616
7	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	LOW	LOW	0,01569
8	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	HIGH	LOW	0,01413
9	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	HIGH	LOW	0,01372
10	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	NORMAL	LOW	0,00714
11	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	NORMAL	LOW	0,00662
12	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	LOW	LOW	0,00618
13	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	NORMAL	LOW	0,00463
14	FALSE	NORMAL	NORMAL	HIGH	HIGH	HIGH	NORMAL	LOW	0,00450
15	FALSE	HIGH	HIGH	HIGH	HIGH	HIGH	NORMAL	LOW	0,00366
16	FALSE	NORMAL	NORMAL	LOW	LOW	LOW	HIGH	LOW	0,00362
17	FALSE	HIGH	HIGH	HIGH	HIGH	HIGH	LOW	LOW	0,00340
18	FALSE	NORMAL	NORMAL	LOW	LOW	LOW	NORMAL	LOW	0,00334
19	FALSE	NORMAL	NORMAL	HIGH	NORMAL	NORMAL	NORMAL	LOW	0,00298
20	FALSE	HIGH	HIGH	HIGH	HIGH	HIGH	HIGH	LOW	0,00297

Table 1 of Figure 2: The 20 most likely value assignments.

[1]Thankfully, it was still possible for us to perform this calculation, despite the fact that the entire list of possible value assignments for this set of variables is almost 45.35 million assignments long.

	MINVOL	FIO2	SAO2	PAP	PRESS	MINVOLSET	CO	BP	probability
1	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	HIGH	HIGH	0.02676
2	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	HIGH	LOW	0.02489
3	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	HIGH	HIGH	0.02175
4	ZERO	NORMAL	LOW	NORMAL	NORMAL	NORMAL	HIGH	HIGH	0.01737
5	ZERO	NORMAL	LOW	NORMAL	LOW	NORMAL	HIGH	HIGH	0.01687
6	ZERO	NORMAL	LOW	NORMAL	NORMAL	NORMAL	HIGH	LOW	0.01616
7	ZERO	NORMAL	LOW	NORMAL	LOW	NORMAL	HIGH	LOW	0.01569
8	ZERO	NORMAL	LOW	NORMAL	NORMAL	NORMAL	HIGH	HIGH	0.01413
9	ZERO	NORMAL	LOW	NORMAL	LOW	NORMAL	HIGH	HIGH	0.01372
10	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	HIGH	NORMAL	0.00714
11	HIGH	NORMAL	HIGH	NORMAL	HIGH	NORMAL	HIGH	HIGH	0.00662
12	HIGH	NORMAL	HIGH	NORMAL	HIGH	NORMAL	HIGH	LOW	0.00618
13	ZERO	NORMAL	LOW	NORMAL	NORMAL	NORMAL	HIGH	NORMAL	0.00463
14	ZERO	NORMAL	LOW	NORMAL	LOW	NORMAL	HIGH	NORMAL	0.00450
15	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	HIGH	HIGH	0.00366
16	HIGH	NORMAL	HIGH	NORMAL	HIGH	NORMAL	NORMAL	HIGH	0.00362
17	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	HIGH	LOW	0.00340
18	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	NORMAL	NORMAL	0.00334
19	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	HIGH	HIGH	0.00298
20	ZERO	NORMAL	LOW	NORMAL	HIGH	NORMAL	HIGH	HIGH	0.00297

Table 2 of Figure 2: The 20 most likely value assignments.

Note that it can happen, that the next-most-likely joint value assignment has more than just a single deviation from the previous-most-likely joint value assignment. Single deviations between rows happen occasionally, but we also see cases where there are many deviations from one row to the next, like between rows 19 and 20, where we see five deviations. In terms of intuition, such differences may seem surprising. However, we could explain them if we would look very closely at the relationships in the network. For a simple example, we also see two deviations between rows 1 and 2, namely the values of TPR and BP. This can be explained easily by the close relationship between TPR and BP in the network, as TPR is a direct parent of BP.

We use the 20 joint value assignments $\mathbf{e}$, as specified in Figure 2, to define 20 different inputs to our implementation. Per input, we will utilise our implementation in order to calculate the different measures of relevance over each variable R of the thirteen intermediate variables $\mathbf{I}$. Thus, we will find 20 results per intermediate variable. We will make observations about them in the next chapter.

Chapter 5

Results

In this chapter, we exclusively share observations about our results. The results themselves are specified in [Appendix B]. They consist of thirteen tables, each of which covers the 20 results relating to one of the intermediate variables. This was hard to integrate nicely with the text, and otherwise just formed a big cluster of tables, interrupting the flow of the paper.

Our observations about the results are listed below. It should be noted that we will sometimes abbreviate intrinsic relevance as IR, and quantified MAP-dependence as QMD. Furthermore, be reminded that MAP-independence is strictly related to quantified MAP-dependence in the following sense: If the quantified measure is 0.00 for some instance, then MAP-independence is a YES for that instance, and otherwise it is a NO. Therefore, we mostly point out comparisons between quantified MAP-dependence and intrinsic relevance, leaving implications about MAP-independence implicit. The only exception to this is the first observation.

- First, we note that within the results for the same variable, different inputs which have a *NO* for MAP-independence sometimes have different quantified MAP-dependence values. This occurs in tables 1, 7, 8, 9 and 10.

- Secondly, looking at results for the same input but across different variables, we observe that the different variables take on more than just two different quantified MAP-dependence values. In none of the 20 inputs do we find that the thirteen variables all take on either value 0.00 or *the same* non-zero value. For instance, in input 1, the set of unique observed QMD values is 0.00, 0.25, 0.33, 0.50. This set includes three different non-zero values.

- Next, we point out that there are four variables for which both quantified MAP-dependence and intrinsic relevance have a value of 0 for all 20 inputs. This can be seen in tables 3, 4, 12 and 13. Both measures indicate that each of these four variables should be considered completely irrelevant to the calculation of MAP, at least for the 20 most likely inputs to the network.

- Continuing, note that there are very few instances where the two measures have the same *non-zero* value. This only happens in three instances found in table 1, so three out off the total 260 instances. There are obviously far more instances where the two measures both have the value zero. This happens in a total of 96 instances, across six variables. Of course, 80 of those instances coincide with the four variables of the previous observation. In either case, out of the total 260 instances, there are 99 instances where the measures have the same value.

- We also see that, in instances where there are differences, it is not always IR that is larger than QMD or vice-versa. Tables 1, 2, 6, 7, 8, 9 and 10 all include instances where the IR-value is higher than the QMD-value, while tables 7, 10 and 11 include instances where the QMD-value is higher than the IR-value.

- Generally, we observe that the value of intrinsic relevance is much more stable across different inputs than the value of quantified MAP-dependence. The value of IR remains the same across all 20 inputs for nine of the thirteen variables. Even in the results per variable where the value of IR does change across inputs, the biggest difference is a change of 0.22 (from 0.69 to 0.91, as observed in table 9). The value of QMD remains the same across all 20 inputs for only five of the thirteen variables. In the results of per variable where the value of QMD does change across inputs, the biggest difference is a change of 0.75 (from 0.00 to 0.75). The only variable for which the value of QMD is stable across all inputs but the value of IR is not, is SHUNT, as seen in table 6.

- Somewhat similar to the previous observation, note that for intrinsic relevance, a table either contains all zeroes, or no zeroes at all. For quantified MAP-dependence, on the other hand, a mix of zeroes and non-zero values happens in tables 5, 7, 9, 10 and 11.

- Finally, we notice that quantified MAP-dependence never has a value higher than 0.75 across all instances. This does happen for intrinsic relevance in tables 8 and 9, even reaching a value of 1.0 in nine of the instances documented in table 8.

Chapter 6

Conclusion

In this paper, we have presented our research about how quantified MAP-dependence compares to the closely related MAP-independence and a different measure of relevance, intrinsic relevance. We made observations about our results in the previous chapter, but left the conclusions to be presented here.

We hypothesised that a quantified version of MAP-independence could be used to identify differences in how relevant each variable in a set of relevant variables is, which we would not be able to do with MAP-independence. We found evidence in support of this when we observed that, across the same input, some of the thirteen intermediate variables had different non-zero values of quantified MAP-dependence. Therefore, we conclude this hypothesis to be true. Furthermore, we noticed that quantified MAP-dependence could also be used to identify differences in how relevant a variable is across different inputs. This is something we had not hypothesised, but it does further motivate our quantification of MAP-independence.

We were also interested in how quantified MAP-dependence compares to intrinsic relevance. Our approach was simply to make a comparison between the two, so we had no hypothesis about this. In our opinion, the evidence we have found indicates that these two measures of relevance seem to describe fundamentally different concepts of what makes an intermediate variable relevant. This may seem obvious, since the measures have different formal definitions. However, it was possible that the two approaches would still lead to very similar results, or even the same results. In such cases, we might have argued that they seem to describe the same concept, despite their different approaches. In our observations, we did see *some* overlap between the measures, especially regarding four variables which both measures regarded as completely irrelevant for all 20 inputs. However, in *the majority* of our results, the two measurements had different values. Furthermore, it seemed like the intrinsic relevance of a variable varied less across different inputs than the quantified MAP-dependence. Likewise, we saw in our results that if a variable has an intrinsic relevance of 0.00 for a single instance, it had an intrinsic relevance of 0.00 for all instances. This was not the same for quantified MAP-dependence. Based on these observations, we conclude that quantified MAP-dependence seems to give more weight to how much impact an intermediate variable has in a specific instance, whereas intrinsic relevance seems to give more of an overall view of how much impact that intermediate variable has in the network as whole.

Keeping in mind that the end goal of applying these measures of relevance is to help justify the decisions made by a decision support system in a specific situation (by identifying which intermediate variables are the most relevant to the calculation of MAP in that situation), we think it makes sense for our measure to capture a lot of variance in how relevant a certain intermediate variable is across different situations. Because we found our results measuring over the 20 most likely evidence inputs to the ALARM network, we think they are meaningful. However, we recognise that we have restricted our research quite a lot. After all, we have only treated the ALARM network, and only measured the results for the 20 most likely inputs. Future research could explore whether our conclusions are supported by evidence based on other Bayesian networks, or whether they hold up when taking a bigger sample from the same network.

An other option for future research relates to the last observation we made about our results. We noticed that quantified MAP-dependence never had a value higher than 0.75 across all our results. This sparked the thought that, perhaps, it is inherently impossible for quantified MAP-dependence, as it was defined in [Chapter 3], to measure a value of 1. Theoretically, we would find a value of 1, if for all possible joint value assignments $\mathbf{r}$, $\text{argmax}_{\mathbf{H}} \Pr(\mathbf{H}, \mathbf{R} = \mathbf{r} \mid \mathbf{e})$ was *not* the same as the optimal solution $\mathbf{h}^*$. However, since the different possible joint value assignments $\mathbf{r}$ have different probabilities, the most likely of them might already *affect* $\mathbf{h}^*$ more than the others. Let us call the most likely value assignment $\mathbf{r}'$. Perhaps it is the case that $\text{argmax}_{\mathbf{H}} \Pr(\mathbf{H}, \mathbf{R} = \mathbf{r}' \mid \mathbf{e}) \equiv \mathbf{h}^*$?

Finally, note that in [9], Kwisthout observed a clear relationship between MAP-independence and the *same-decision probability* defined by Choi and colleagues in [14]. It would be interesting to see future research exploring whether a similarly clear relationship could be found between quantified MAP-dependence and the same-decision probability, and how the two measures compare in general.

To summarise: We defined a new measure, quantified MAP-dependence. We made the first implementation of both MAP-independence and quantified MAP-dependence, and also implemented intrinsic relevance. In our research, we applied these three measures to the ALARM network, and made comparisons between them based on the results. We concluded that quantified MAP-dependence improves upon MAP-independence by allowing us to identify different levels of relevance, and also that quantified MAP-dependence seems to describe a fundamentally different idea of what makes a variable relevant, compared to intrinsic relevance. However, we recognise that we have left much room for further research into the new measurement.

Bibliography

[1] E. Kyrimi, K. Dube, N. Fenton, *et al.*, "Bayesian networks in healthcare: What is preventing their adoption?" *Artificial Intelligence in Medicine*, vol. 116, p. 102 079, 2021.

[2] A. S. Cofino, R. Cano, C. Sordo, and J. M. Gutierrez, "Bayesian networks for probabilistic weather prediction," in *15th Eureopean Conference on Artificial Intelligence (ECAI)*, Citeseer, 2002.

[3] C. S. Vlek, H. Prakken, S. Renooij, and B. Verheij, "A method for explaining Bayesian networks for legal evidence with scenarios," *Artificial Intelligence and Law*, vol. 24, no. 3, pp. 285–324, 2016.

[4] D. Gunning, M. Stefik, J. Choi, T. Miller, S. Stumpf, and G.-Z. Yang, "XAI—Explainable artificial intelligence," *Science Robotics*, vol. 4, no. 37, eaay7120, 2019.

[5] R. Confalonieri, L. Coba, B. Wagner, and T. R. Besold, "A historical perspective of explainable artificial intelligence," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 11, no. 1, e1391, 2021.

[6] M. El-Assady, W. Jentner, R. Kehlbeck, *et al.*, "Towards XAI: Structuring the processes of explanations," in *Proceedings of the ACM Workshop on Human-Centered Machine Learning, Glasgow, UK*, vol. 4, 2019.

[7] D. Doran, S. Schulz, and T. R. Besold, "What does explainable AI really mean? A new conceptualization of perspectives," *arXiv preprint arXiv:1710.00794*, 2017.

[8] X. Wang and M. Yin, "Are explanations helpful? A comparative study of the effects of explanations in AI-assisted decision-making," in *26th International Conference on Intelligent User Interfaces*, 2021, pp. 318–328.

[9] J. Kwisthout, "Explainable AI using MAP-Independence," in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, J. Vejnarová and N. Wilson, Eds., Cham: Springer International Publishing, 2021, pp. 243–254, ISBN: 978-3-030-86772-0.

[10] I. A. Beinlich, H. J. Suermondt, R. M. Chavez, and G. F. Cooper, "The ALARM monitoring system: A case study with two probabilistic inference techniques for belief networks," in *AIME 89*, Springer, 1989, pp. 247–256.

[11] J. Kwisthout, "Most frugal explanations in Bayesian networks," *Artificial Intelligence*, vol. 218, pp. 56–73, 2015. DOI: `http://dx.doi.org/10.1016/j.artint.2014.10.001`.

[12] A. Darwiche, *Modeling and reasoning with Bayesian networks*. Cambridge university press, 2009.

[13] M. J. Druzdzel and H. J. Suermondt, "Relevance in probabilistic models: "backyards" in a "small world"," in *Working notes of the AAAI–1994 Fall Symposium Series: Relevance*, 1994, pp. 60–63.

[14] A. Choi, Y. Xue, and A. Darwiche, "Same-decision probability: A confidence measure for threshold-based decisions," *International Journal of Approximate Reasoning*, vol. 53, no. 9, pp. 1415–1428, 2012.

Appendix A

Implementation

The complete implementation we have used can be found on github, using the following link: `https://github.com/MvanElteren/bachelorThesis-project`

Appendix B

Results

In this appendix, we specify the results. This is done in the form of thirteen tables, each of which covers the 20 results relating to one of the intermediate variables.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.33	NO	0.67
2	0.33	NO	0.67
3	0.33	NO	0.67
4	0.33	NO	0.67
5	0.33	NO	0.67
6	0.33	NO	0.67
7	0.33	NO	0.67
8	0.33	NO	0.67
9	0.33	NO	0.67
10	0.33	NO	0.67
11	0.33	NO	0.67
12	0.33	NO	0.67
13	0.33	NO	0.67
14	0.33	NO	0.67
15	0.67	NO	0.67
16	0.33	NO	0.67
17	0.67	NO	0.67
18	0.33	NO	0.67
19	0.33	NO	0.67
20	0.67	NO	0.67

Table 1: R = LVEDVOLUME.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.00	YES	0.67
2	0.00	YES	0.67
3	0.00	YES	0.67
4	0.00	YES	0.67
5	0.00	YES	0.67
6	0.00	YES	0.67
7	0.00	YES	0.67
8	0.00	YES	0.67
9	0.00	YES	0.67
10	0.00	YES	0.67
11	0.00	YES	0.67
12	0.00	YES	0.67
13	0.00	YES	0.67
14	0.00	YES	0.67
15	0.33	NO	0.67
16	0.00	YES	0.67
17	0.33	NO	0.67
18	0.00	YES	0.67
19	0.00	YES	0.67
20	0.33	NO	0.67

Table 2: R = STROKEVOLUME.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.00	YES	0.00
2	0.00	YES	0.00
3	0.00	YES	0.00
4	0.00	YES	0.00
5	0.00	YES	0.00
6	0.00	YES	0.00
7	0.00	YES	0.00
8	0.00	YES	0.00
9	0.00	YES	0.00
10	0.00	YES	0.00
11	0.00	YES	0.00
12	0.00	YES	0.00
13	0.00	YES	0.00
14	0.00	YES	0.00
15	0.00	YES	0.00
16	0.00	YES	0.00
17	0.00	YES	0.00
18	0.00	YES	0.00
19	0.00	YES	0.00
20	0.00	YES	0.00

Table 3: R = ERRLOWOUTPUT.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.00	YES	0.00
2	0.00	YES	0.00
3	0.00	YES	0.00
4	0.00	YES	0.00
5	0.00	YES	0.00
6	0.00	YES	0.00
7	0.00	YES	0.00
8	0.00	YES	0.00
9	0.00	YES	0.00
10	0.00	YES	0.00
11	0.00	YES	0.00
12	0.00	YES	0.00
13	0.00	YES	0.00
14	0.00	YES	0.00
15	0.00	YES	0.00
16	0.00	YES	0.00
17	0.00	YES	0.00
18	0.00	YES	0.00
19	0.00	YES	0.00
20	0.00	YES	0.00

Table 4: R = ERRCAUTER.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.33	NO	0.00
2	0.33	NO	0.00
3	0.33	NO	0.00
4	0.00	YES	0.00
5	0.00	YES	0.00
6	0.00	YES	0.00
7	0.00	YES	0.00
8	0.00	YES	0.00
9	0.00	YES	0.00
10	0.33	NO	0.00
11	0.33	NO	0.00
12	0.33	NO	0.00
13	0.00	YES	0.00
14	0.00	YES	0.00
15	0.33	NO	0.00
16	0.33	NO	0.00
17	0.33	NO	0.00
18	0.33	NO	0.00
19	0.33	NO	0.00
20	0.33	NO	0.00

Table 5: R = PVSAT.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.00	YES	0.28
2	0.00	YES	0.28
3	0.00	YES	0.28
4	0.00	YES	0.34
5	0.00	YES	0.41
6	0.00	YES	0.34
7	0.00	YES	0.41
8	0.00	YES	0.34
9	0.00	YES	0.41
10	0.00	YES	0.28
11	0.00	YES	0.30
12	0.00	YES	0.30
13	0.00	YES	0.34
14	0.00	YES	0.41
15	0.00	YES	0.28
16	0.00	YES	0.30
17	0.00	YES	0.28
18	0.00	YES	0.28
19	0.00	YES	0.28
20	0.00	YES	0.28

Table 6: R = SHUNT.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.00	YES	0.50
2	0.00	YES	0.50
3	0.00	YES	0.50
4	0.00	YES	0.50
5	0.25	NO	0.50
6	0.00	YES	0.50
7	0.25	NO	0.50
8	0.00	YES	0.50
9	0.25	NO	0.50
10	0.00	YES	0.50
11	0.75	NO	0.50
12	0.75	NO	0.50
13	0.00	YES	0.50
14	0.25	NO	0.50
15	0.00	YES	0.50
16	0.75	NO	0.50
17	0.00	YES	0.50
18	0.00	YES	0.50
19	0.00	YES	0.50
20	0.00	YES	0.50

Table 7: R = VENTMACH.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.50	NO	1.00
2	0.50	NO	1.00
3	0.50	NO	1.00
4	0.75	NO	0.91
5	0.75	NO	0.91
6	0.75	NO	0.91
7	0.75	NO	0.91
8	0.75	NO	0.91
9	0.75	NO	0.91
10	0.50	NO	1.00
11	0.75	NO	0.92
12	0.75	NO	0.92
13	0.75	NO	0.91
14	0.75	NO	0.91
15	0.50	NO	1.00
16	0.75	NO	0.92
17	0.50	NO	1.00
18	0.50	NO	1.00
19	0.50	NO	1.00
20	0.50	NO	1.00

Table 8: R = VENTTUBE.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.25	NO	0.69
2	0.25	NO	0.69
3	0.25	NO	0.69
4	0.25	NO	0.91
5	0.00	YES	0.84
6	0.25	NO	0.91
7	0.00	YES	0.84
8	0.25	NO	0.91
9	0.00	YES	0.84
10	0.25	NO	0.69
11	0.75	NO	0.78
12	0.75	NO	0.78
13	0.25	NO	0.91
14	0.00	YES	0.84
15	0.25	NO	0.69
16	0.75	NO	0.78
17	0.25	NO	0.69
18	0.25	NO	0.69
19	0.25	NO	0.69
20	0.25	NO	0.69

Table 9: R = VENTLUNG.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.25	NO	0.47
2	0.25	NO	0.47
3	0.25	NO	0.47
4	0.00	YES	0.50
5	0.00	YES	0.50
6	0.00	YES	0.50
7	0.00	YES	0.50
8	0.00	YES	0.50
9	0.00	YES	0.50
10	0.25	NO	0.47
11	0.75	NO	0.50
12	0.75	NO	0.50
13	0.00	YES	0.50
14	0.00	YES	0.50
15	0.25	NO	0.47
16	0.75	NO	0.50
17	0.25	NO	0.47
18	0.25	NO	0.47
19	0.25	NO	0.47
20	0.25	NO	0.47

Table 10: R = VENTALV.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.33	NO	0.00
2	0.33	NO	0.00
3	0.33	NO	0.00
4	0.00	YES	0.00
5	0.00	YES	0.00
6	0.00	YES	0.00
7	0.00	YES	0.00
8	0.00	YES	0.00
9	0.00	YES	0.00
10	0.33	NO	0.00
11	0.33	NO	0.00
12	0.33	NO	0.00
13	0.00	YES	0.00
14	0.00	YES	0.00
15	0.33	NO	0.00
16	0.33	NO	0.00
17	0.33	NO	0.00
18	0.33	NO	0.00
19	0.33	NO	0.00
20	0.33	NO	0.00

Table 11: R = ARTCO2.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.00	YES	0.00
2	0.00	YES	0.00
3	0.00	YES	0.00
4	0.00	YES	0.00
5	0.00	YES	0.00
6	0.00	YES	0.00
7	0.00	YES	0.00
8	0.00	YES	0.00
9	0.00	YES	0.00
10	0.00	YES	0.00
11	0.00	YES	0.00
12	0.00	YES	0.00
13	0.00	YES	0.00
14	0.00	YES	0.00
15	0.00	YES	0.00
16	0.00	YES	0.00
17	0.00	YES	0.00
18	0.00	YES	0.00
19	0.00	YES	0.00
20	0.00	YES	0.00

Table 12: R = CATECHOL.

input nr.	quantified MAP-dependence	MAP-independence	intrinsic relevance
1	0.00	YES	0.00
2	0.00	YES	0.00
3	0.00	YES	0.00
4	0.00	YES	0.00
5	0.00	YES	0.00
6	0.00	YES	0.00
7	0.00	YES	0.00
8	0.00	YES	0.00
9	0.00	YES	0.00
10	0.00	YES	0.00
11	0.00	YES	0.00
12	0.00	YES	0.00
13	0.00	YES	0.00
14	0.00	YES	0.00
15	0.00	YES	0.00
16	0.00	YES	0.00
17	0.00	YES	0.00
18	0.00	YES	0.00
19	0.00	YES	0.00
20	0.00	YES	0.00

Table 13: R = HR.