

The YODO algorithm: An efficient computational framework for sensitivity analysis in Bayesian networks

Rafael Ballester-Ripoll, Manuele Leonelli

^aSchool of Science and Technology, IE University, Madrid, Spain

Abstract

Sensitivity analysis measures the influence of a Bayesian network’s parameters on a quantity of interest defined by the network, such as the probability of a variable taking a specific value. In the literature, this influence is often measured by computing the partial derivative with respect to the network parameters. However, this can become computationally very expensive in large networks with thousands of parameters. We propose an algorithm combining automatic differentiation and exact inference to efficiently calculate the sensitivity measures in a single pass. It first marginalizes the whole network once, using e.g. variable elimination, and then backpropagates this operation to obtain the gradient with respect to all input parameters. Our method can be used for one-way and multi-way sensitivity analysis and the derivation of admissible regions. Simulation studies highlight the efficiency of our algorithm by scaling it to massive networks with up to 100’000 parameters and investigate the feasibility of generic multi-way analyses. Our routines are also showcased over two medium-sized Bayesian networks: the first modeling the country-risks of a humanitarian crisis, the second studying the relationship between the use of technology and the psychological effects of forced social isolation during the COVID-19 pandemic. An implementation of the methods using the popular machine learning library PyTorch is freely available.

Keywords: Automatic differentiation, Bayesian networks, COVID-19, PyTorch, Sensitivity analysis

1. Introduction

Probabilistic graphical models, and specifically Bayesian networks (BNs), are a class of models that are widely used for risk assessment of complex operational systems in a variety of domains. The main reason for their success is that they provide an efficient and intuitive framework to represent the joint probability of a

vector of variables of interest using a simple graph. Their use to assess the reliability of engineering, medical and ecological systems, among many others, is becoming increasingly popular. Sensitivity analysis is a critical step for any applied real-world analysis to assess the importance of various risk factors and to evaluate the overall safety of the system under study (see e.g. Goerlandt and Islam, 2021; Makaba et al., 2021; Zio et al., 2022, for some recent examples).

As noticed by Rohmer (2020), sensitivity analysis in BNs is usually *local*, in the sense that it measures the effect of a small number of parameter variations on output probabilities of interest, while other parameters are kept fixed. In the case of a single parameter variation, sensitivity analysis is usually referred to as *one-way*; otherwise, when more than one parameter is varied, it is called *multi-way*. Although recently there has been an increasing interest in proposing *global* sensitivity methods for BNs measuring how different factors *jointly* influence some function of the model’s output (see e.g. Ballester-Ripoll and Leonelli, 2022a; Li and Mahadevan, 2018), the focus of this paper still lies in local sensitivity methods.

Local sensitivity analysis in BNs can be broken down into two main steps. First, some parameters of the model are varied, and the effect of these variations on output probabilities of interest is investigated. For this purpose, a simple mathematical function, usually termed *sensitivity function*, describes an output probability of interest as a function of the BN parameters (Castillo et al., 1997; Coupé and van der Gaag, 2002). Furthermore, some specific properties of such a function can be computed, for instance, the *sensitivity value* or the *vertex proximity*, which give an overview of how sensitive the probability of interest is to variations of the associated parameter (van der Gaag et al., 2007). Second, once parameter variations are identified, their effect is summarized by a distance or divergence measure between the original and the varied distributions underlying the BN, most commonly the Chan-Darwiche distance (Chan and Darwiche, 2005) or the well-known Kullback-Leibler divergence.

As demonstrated by Kwisthout and van der Gaag (2008), the derivation of both the sensitivity function and its associated properties is computationally very demanding. In Ballester-Ripoll and Leonelli (2022b), we introduced a novel, computationally highly-efficient method to compute all sensitivity measures of interest in one-way sensitivity analysis, which takes advantage of backpropagation and is easy to compute thanks to automatic differentiation. We now also demonstrate how the algorithm can be utilized for more generic multi-way sensitivity analyses and for deriving admissible regions (van der Gaag and Renooij, 2001). Simulation studies show the efficiency of the approach by processing massive networks in a few seconds and demonstrate when multi-way analyses are computationally feasible. Two practical applications from real-world datasets further showcase the insights sensitivity

measures can provide and the efficiency of the implemented routines.

We have open-sourced a Python implementation using the popular machine learning library PyTorch¹, contributing to the recent effort of promoting sensitivity analysis (Douglas-Smith et al., 2020).

2. Bayesian networks and sensitivity analysis

A BN is a probabilistic graphical model defining a factorization of the probability mass function (PMF) of a random vector using a directed acyclic graph (DAG) (Darwiche, 2009a; Pearl, 1988). More formally, let $[p] = \{1, \dots, p\}$ and $\mathbf{Y} = (Y_i)_{i \in [p]}$ be a random vector of interest with sample space $\mathbb{Y} = \times_{i \in [p]} \mathbb{Y}_i$. A BN defines the PMF $P(\mathbf{Y} = \mathbf{y})$, for $\mathbf{y} \in \mathbb{Y}$, as a product of simpler conditional PMFs as follows:

$$P(\mathbf{Y} = \mathbf{y}) = \prod_{i \in [p]} P(Y_i = y_i \mid \mathbf{Y}_{\Pi_i} = \mathbf{y}_{\Pi_i}), \quad (1)$$

where \mathbf{Y}_{Π_i} are the parents of Y_i in the DAG associated to the BN.

The definition of the PMF over \mathbf{Y} , which would require defining an exponential number of probabilities, is thus simplified in terms of one-dimensional conditional PMFs. The coefficients of these functions are henceforth referred to as the parameters $\boldsymbol{\theta}$ of the model. The DAG structure may be either elicited by experts or learned from data using structural learning algorithms. Similarly, the associated parameters $\boldsymbol{\theta}$ can be either elicited or learned using frequentist or Bayesian approaches. No matter the method used, we assume that a value for these parameters $\boldsymbol{\theta}$ has been chosen, which we refer to as the *original value* $\boldsymbol{\theta}^0$.

The DAG associated with a BN provides an intuitive overview of the relationships between variables of interest. However, it does also provide a framework to assess if any generic conditional independence holds for a specific subset of the variables via the so-called d-separation criterion (see e.g. Pearl, 1988). Furthermore, the DAG provides a framework for the efficient propagation of probabilities and evidence via algorithms that take advantage of the structure of the underlying DAG.

2.1. One-way sensitivity analysis

In practical applications, it is fundamental to extensively assess the implications of the chosen parameter values $\boldsymbol{\theta}^0$ to outputs of the model. In the context of BNs, this study is usually referred to as *sensitivity analysis*, which can be further used

¹Available at <https://github.com/rballester/yodo>.

during the model-building process as showcased by Coupé et al. (2000). Let Y_O be an output variable of interest and \mathbf{Y}_E be *evidential* variables, those that may be observed. The interest is in then studying how $P(Y_O = y_O \mid \mathbf{Y}_E = \mathbf{y}_E)$ varies when a parameter θ_i is varied. In particular, $P(Y_O = y_O \mid \mathbf{Y}_E = \mathbf{y}_E)$ seen as a function of θ_i is called *sensitivity function* and denoted as $f(\theta_i)$.

2.2. Proportional covariation

Notice that when an input θ_i is varied from its original value θ_i^0 , the parameters from the same conditional PMF need to *covary* to respect the sum-to-one condition of probabilities. When variables are binary, this is automatic since one parameter must be equal to one minus the other. However, for variables taking more than two levels, this covariation can be done in several ways (Renooij, 2014). We henceforth assume that whenever a parameter is varied from its original value θ_i^0 to a new value θ_i , then every parameter θ_j from the same conditional PMF is *proportionally covaried* (Laskey, 1995) from its original value θ_j^0 :

$$\theta_j(\theta_i) = \frac{1 - \theta_i}{1 - \theta_i^0} \theta_j^0. \quad (2)$$

Proportional covariation has been studied extensively, and its choice is motivated by a wide array of theoretical properties (Chan and Darwiche, 2005; Leonelli et al., 2017; Leonelli and Riccomagno, 2022; Renooij, 2014).

Under the assumption of proportional covariation, Castillo et al. (1997) and Coupé and van der Gaag (2002) demonstrated that the sensitivity function is the ratio of two linear functions:

$$f(\theta_i) = \frac{c_0 + c_i \theta_i}{d_0 + d_i \theta_i}, \quad (3)$$

where $c_0, c_i, d_0, d_i \in \mathbb{R}_+$. van der Gaag et al. (2007) noticed that the above expression coincides with the fragment of a rectangular hyperbola, which can be generally written as

$$f(\theta_i) = \frac{r}{\theta_i - s} + t, \quad (4)$$

where

$$s = -\frac{d_0}{d_i}, \quad t = \frac{c_i}{d_i}, \quad r = \frac{c_0}{d_i} + st. \quad (5)$$

2.2.1. Sensitivity values

The *sensitivity value* describes the effect of infinitesimally small shifts in the parameter's original value on the probability of interest and is defined as the absolute value of the first derivative of the sensitivity function at the original value of the

parameter, i.e. $|f'(\theta_i^0)|$. This can be found by simply differentiating the sensitivity function as

$$|f'(\theta_i^0)| = \frac{|c_i d_0 - c_0 d_i|}{(d_i \theta_i^0 + d_0)^2}. \quad (6)$$

The higher the sensitivity value, the more sensitive the output probability to small changes in the parameter's original value. As a rule of thumb, parameters having a sensitivity value larger than one may require further investigation.

Notice that when \mathbf{Y}_E is empty, i.e. the output probability of interest is marginal, the sensitivity function is linear in θ_i . The sensitivity value is the same regardless of the original θ_i^0 . Therefore, in this case, the absolute value of the gradient is sufficient to quantify the effect of a parameter on an output probability of interest.

2.2.2. Vertex proximity

van der Gaag et al. (2007) further noticed that parameters for which the sensitivity value is small may still be such that the conditional output probability of interest is very sensitive to their variations. This happens when the original parameter value is close to the *vertex* of the sensitivity function, defined as the point θ_i^v at which the sensitivity value is equal to one, i.e.

$$|f'(\theta_i^v)| = 1. \quad (7)$$

The vertex can be derived from the equation of the sensitivity function as

$$\theta_i^v = \begin{cases} s + \sqrt{|r|}, & \text{if } s < 0, \\ s - \sqrt{|r|}, & \text{if } s > 0. \end{cases} \quad (8)$$

Notice that the case $s = 0$ is not contemplated since it would coincide with a linear sensitivity function, not a hyperbolic one.

Vertex proximity is defined as the absolute difference $|\theta_i^0 - \theta_i^v|$. The smaller the vertex proximity, the more sensitive the output probabilities may be to parameter variations, even when the sensitivity value is small.

2.2.3. Other metrics

Given the coefficients c_0, c_i, d_0, d_i of Equation (3), it is straightforward to derive any property of the sensitivity function besides the sensitivity value and the vertex proximity. Here we propose the use of two additional metrics. The first is the absolute value of the second derivative of the sensitivity function at the original parameter value, which can be easily computed as:

$$|f''(\theta_i^0)| = \frac{2d_i |c_i d_0 - c_0 d_i|}{(d_i \theta_i^0 + d_0)^3}. \quad (9)$$

Similarly to the sensitivity value, high values of the second derivative at θ_i^0 indicate parameters that could highly impact the probability of interest.

The second measure is the maximum of the first derivative of the sensitivity function over the interval $[0, 1]$ in absolute value, which we find easily by noting that the denominator of Equation (6) is a parabola:

$$\max_{\theta_i \in [0,1]} |f'(\theta_i)| = \begin{cases} \infty & \text{if } -d_0/d_i \in [0, 1] \\ \max\{|c_i d_0 - c_0 d_i|/d_0^2, |c_i d_0 - c_0 d_i|/(d_i + d_0)^2\} & \text{otherwise.} \end{cases} \quad (10)$$

Again high values indicate parameters whose variations can lead to a significant change in the output probability of interest.

2.3. Multi-way sensitivity analysis

In many practical applications, there is interest in assessing the effect of simultaneous variations of multiple parameters on the output of interest. This is called a *multi-way sensitivity analysis*. Although there have been some attempts to study the theoretical properties and computational efficiency of these more generic analyses (see e.g. Bolt and Renooij, 2014; Chan and Darwiche, 2004; Kjaerulff and van der Gaag, 2000; Leonelli et al., 2017; Leonelli and Riccomagno, 2022), in practice, they are not as common as one-way analyses.

2.3.1. General formulation

Suppose now that n parameters $\boldsymbol{\theta}_n = (\theta_1, \dots, \theta_n)$ are simultaneously varied. By default, these parameters are taken from different conditional PMFs so that they are independent of each other (van der Gaag et al., 2007). In the binary case, this is natural since only one parameter per PMF can be varied since the other is functionally related. The other parameters from conditional PMFs including $\theta_1, \dots, \theta_n$, are proportionally covaried, as for the one-way analysis (see Leonelli and Riccomagno, 2022, for a formal discussion). The effect of varying the parameters $\boldsymbol{\theta}_n$ on a probability of interest $P(Y_O = y_O \mid \mathbf{Y}_E = \mathbf{y}_E)$ is captured by the n-way sensitivity function, which is equal to

$$f(\boldsymbol{\theta}_n) = \frac{\sum_{K \in \mathcal{P}([n])} c_K \prod_{i \in K} \theta_i}{\sum_{K \in \mathcal{P}([n])} d_K \prod_{i \in K} \theta_i}, \quad (11)$$

where \mathcal{P} denotes the power set and $c_K, d_K \in \mathbb{R}$, $K \in \mathcal{P}([n])$, are constants computed from the non-varied parameters. For instance, a 2-way sensitivity function can be written as:

$$f(\theta_1, \theta_2) = \frac{c_0 + c_1\theta_1 + c_2\theta_2 + c_{12}\theta_1\theta_2}{d_0 + d_1\theta_1 + d_2\theta_2 + d_{12}\theta_1\theta_2} \quad (12)$$

An n-way sensitivity function, in general, requires the computation of 2^{n+1} constants and is thus computationally expensive. Furthermore, the number of combinations of parameters for which the sensitivity function has to be constructed increases: see Section 3.2 for a discussion.

2.3.2. Maximum n-way sensitivity values

While for one-way sensitivity analysis, one can uniquely talk about the derivative of the sensitivity function, for multi-valued functions, there are multiple directions at which the derivative could be computed, as noted by (Bolt and Renooij, 2014), and hence the notion of *directional derivative*. However, basic calculus tells us that the maximum directional derivative of a function f at a point θ_n equals the length of the gradient vector at θ_n , i.e. $|\nabla f(\theta_n)|$. This observation led to the definition of the sensitivity value for an n-way sensitivity function as the maximum one out of all possible directional derivatives (Bolt and Renooij, 2014). For a vector of parameters θ_n with original values θ_n^0 the *maximum n-way sensitivity value* is defined as

$$sv_{\max}^{\theta_n} = |\nabla f(\theta_n^0)|, \quad (13)$$

where f is the associated n-way sensitivity function.

By definition, the maximum n-way sensitivity value would first require the derivation of the n-way sensitivity function and, subsequently, the computation of its gradient. As noticed already, this direct approach would be computationally too expensive. However, Bolt and Renooij (2014) demonstrated that $sv_{\max}^{\theta_n}$ could be easily computed from the sensitivity values of one-way sensitivity functions. Let c_0^i, c_i, d_0^i, d_i be the coefficients of the one-way sensitivity function for the variation of the parameter θ_i in θ_n . Then:

$$sv_{\max}^{\theta_n} = \frac{1}{P(\mathbf{Y}_E = \mathbf{y}_E)^2} \sqrt{\sum_{i \in [n]} (c_i d_0^i - c_0^i d_i)^2}. \quad (14)$$

Therefore if an efficient method for computing the coefficients of one-way sensitivity functions exists, then maximum n-way sensitivity values can be equally efficiently derived.

2.4. Admissible regions

In many applied situations, the object of interest is not a probability per se, but rather the most likely value of a variable, possibly conditional on a specific subset of evidence. This is the case for classification problems where a Bayes classifier is used: an unlabeled observation exhibiting a specific evidence pattern is classified according to the most likely value. BNs designed explicitly for this task are usually called Bayesian network classifiers (Bielza and Larranaga, 2014; Friedman et al., 1997).

Although sensitivity methods for this type of classification problem have been discussed (Bolt and van der Gaag, 2017), sensitivity values and related measures are sometimes not particularly useful. van der Gaag and Renooij (2001) demonstrated that parameters may have a small sensitivity value, yet change the most likely value of Y_O when slightly altered. For this reason, they introduced the concept of *admissible region*, which captures the extent to which a parameter can be varied without inducing a change in the most likely value for the variable of interest.

For ease of notation, we consider here a variable of interest Y_O taking two possible levels y_O and y'_O (thus, we consider the most common binary classification problem). Consider also possible evidence $\mathbf{Y}_E = \mathbf{y}_E$, a perturbed parameter θ_i and suppose that $P(Y_O = y_O | \mathbf{Y}_E = \mathbf{y}_E) > P(Y_O = y'_O | \mathbf{Y}_E = \mathbf{y}_E)$, without loss of generality. The admissible region R_i is formally defined as the interval of values for θ_i

$$(\max\{\theta_i^0 - r, 0\}, \min\{\theta_i^0 + s, 1\}), \quad r, s, \in \mathbb{R}, \quad (15)$$

for which

$$P(Y_O = y_O | \mathbf{Y}_E = \mathbf{y}_E) > P(Y_O = y'_O | \mathbf{Y}_E = \mathbf{y}_E). \quad (16)$$

The wider the interval R_i , the less influential the parameter is for the most likely value.

van der Gaag and Renooij (2001) and van der Gaag et al. (2007) already demonstrated that such regions could be computed from the one-way sensitivity functions by identifying the points at which the sensitivity functions intersect. However, they did not explicitly write the admissible regions as a function of the sensitivity functions' coefficients to our knowledge. Let c_0, c_i, d_0, d_i be the coefficients of the sensitivity function for the event $y_O | \mathbf{y}_E$. It follows that the sensitivity function for $y'_O | \mathbf{y}_E$ must be equal to

$$\frac{(d_0 - c_0) + (d_i - c_i)\theta_i}{d_0 + d_i\theta_i}. \quad (17)$$

By equating the two sides of Eq. 16, we find that

$$R_i = \begin{cases} \left(0, \min\left\{\frac{d_0-2c_0}{2c_i-d_i}, 1\right\}\right) & \text{if } \theta_i^0 \leq \frac{d_0-2c_0}{2c_i-d_i} \\ \left(\max\left\{0, \frac{d_0-2c_0}{2c_i-d_i}\right\}, 1\right) & \text{otherwise} \end{cases} \quad (18)$$

In the case of $E = \emptyset$, i.e. no evidence, the expression for the admissible regions simplifies to:

$$R_i = \begin{cases} \left(0, \min\left\{\frac{1-2c_0}{2c_i}, 1\right\}\right) & \text{if } \theta_i^0 \leq \frac{1-2c_0}{2c_i} \\ \left(\max\left\{0, \frac{1-2c_0}{2c_i}\right\}, 1\right) & \text{otherwise} \end{cases} \quad (19)$$

Therefore, given an efficient method to compute one-way sensitivity functions, admissible regions for all individual parameters can be equally efficiently derived.

3. The YODO method

The YODO (You Only Derive Once) method was first introduced in Ballester-Ripoll and Leonelli (2022b) to compute the one-way sensitivity measures discussed in Sections 2.2.1-2.2.3. We first review it and then discuss its use in multi-way sensitivity analysis.

3.1. YODO for one-way sensitivity analysis

3.1.1. First case: Marginal probability as a function of interest

Suppose $f(\theta_i) = P(Y_O = y_O) = c_0 + c_i\theta_i$ assuming proportional covariation as θ_i varies. Let $\theta_{j_1}, \dots, \theta_{j_n}$ be the other parameters of the same conditional PMF as θ_i , i.e. they are all bound by the sum-to-one constraint $\theta_i + \theta_{j_1} + \dots + \theta_{j_n} = 1$. First, we rewrite f as

$$f(\theta_i) = g(\theta_i, \theta_{j_1}(\theta_i), \dots, \theta_{j_n}(\theta_i)) \quad (20)$$

and we show how to obtain $f'(\theta_i)$ provided that we can compute the gradient ∇g with respect to symbols $\theta_i, \theta_{j_1}, \dots, \theta_{j_n}$ (see Section 3.1.3 for details on the latter).

By the generalized chain rule, it holds that

$$f'(\theta_i) = \frac{\partial g}{\partial \theta_i} \cdot 1 + \frac{\partial g}{\partial \theta_{j_1}} \cdot \frac{d\theta_{j_1}}{d\theta_i} + \dots + \frac{\partial g}{\partial \theta_{j_n}} \cdot \frac{d\theta_{j_n}}{d\theta_i}. \quad (21)$$

By deriving Equation (2), we have that for all $1 \leq m \leq n$:

$$\frac{d\theta_{j_m}}{d\theta_i} = \frac{-\theta_{j_m}^0}{1 - \theta_i^0} \quad (22)$$

and, therefore,

$$f'(\theta_i) = \frac{\partial g}{\partial \theta_i} - \frac{(\partial g / \partial \theta_{j_1}) \cdot \theta_{j_1}^0 + \dots + (\partial g / \partial \theta_{j_n}) \cdot \theta_{j_n}^0}{1 - \theta_i^0}. \quad (23)$$

Last, since $f(\theta_i) = P(\mathbf{Y}_O = \mathbf{y}_O) = c_0 + c_i \theta_i$, we easily find the parameters c_0, c_i :

$$\begin{cases} c_i = f'(\theta_i^0) \\ c_0 = P(\mathbf{Y}_O = \mathbf{y}_O) - c_i \theta_i^0. \end{cases} \quad (24)$$

3.1.2. Second case: Conditional probability as a function of interest

When $f(\theta_i) = P(Y_O = y_O \mid \mathbf{Y}_E = \mathbf{y}_E) = P(Y_O = y_O, \mathbf{Y}_E = \mathbf{y}_E) / P(\mathbf{Y}_E = \mathbf{y}_E)$, we simply repeat the procedure from Sec. 3.1.1 twice:

1. We first apply it to $P(Y_O = y_O, \mathbf{Y}_E = \mathbf{y}_E)$ to obtain c_0 and c_i ;
2. we then apply it to $P(\mathbf{Y}_E = \mathbf{y}_E)$ to obtain d_0 and d_i .

3.1.3. Computing the gradient ∇g

Let $\mathbf{Y}_K = \mathbf{y}_K$ be a subset of the network variables taking some evidence values (this could be $K = O$ or $K = O \cup E$; hence we cover the two cases above).

We start by moralizing the BN into a Markov random field (MRF) \mathcal{M} . This marries all variable parents together and, for each conditional probability table (now called *potential*), drops the sum-to-one constraint; see e.g. (Darwiche, 2009b) for more details. Next, we impose the evidence $\mathbf{Y}_K = \mathbf{y}_K$ by defining $\mathcal{M}^{\mathbf{Y}_K = \mathbf{y}_K}$ as a new MRF that results from substituting each potential $\Phi_{i_1, \dots, i_M}(x_{i_1}, \dots, x_{i_M})$ by a new potential $\widehat{\Phi}_{i_1, \dots, i_M}$ defined as follows:

$$\widehat{\Phi}_{i_1, \dots, i_M}(Y_{i_1} = x_{i_1}, \dots, Y_{i_M} = x_{i_M}) = \begin{cases} 0 & \text{if } \exists m, k \mid i_m = k \wedge x_{i_m} \neq y_{i_m} \\ \Phi_{i_1, \dots, i_M}(Y_{i_1} = x_{i_1}, \dots, Y_{i_M} = x_{i_M}) & \text{otherwise} \end{cases} \quad (25)$$

In other words, we copy the original potential but zero-out all entries that do not honor the assignment of values $\mathbf{Y}_K = \mathbf{y}_K$. See Table 1 for an example using a bivariate potential.

Intuitively, the modified MRF $\mathcal{M}^{\mathbf{Y}_K = \mathbf{y}_K}$ represents the unnormalized probability for all variable assignments that are compatible with $\mathbf{Y}_K = \mathbf{y}_K$. In particular, if $\mathcal{M}_{\mathbf{Y}_K}$ denotes the marginalization of a network \mathcal{M} over all variables in \mathbf{Y}_K , we have that $(\mathcal{M}^{\mathbf{Y}_K = \mathbf{y}_K})_{\mathbf{Y}} = P(\mathbf{Y}_K = \mathbf{y}_K)$. In other words, computing g reduces to marginalizing

	$Y_2 = 1$	$Y_2 = 2$	$Y_2 = 3$
$Y_1 = 1$	0.8	0.1	0.1
$Y_1 = 2$	0.3	0.5	0.2
$Y_1 = 3$	0.1	0.2	0.7

(a) $\Phi_{1,2}(y_1, y_2)$

	$Y_2 = 1$	$Y_2 = 2$	$Y_2 = 3$
$Y_1 = 1$	0.0	0.0	0.1
$Y_1 = 2$	0.0	0.0	0.2
$Y_1 = 3$	0.0	0.0	0.7

(b) $\hat{\Phi}_{1,2}(y_1, y_2)$

Table 1: Left: example potential of an MRF \mathcal{M} for variables Y_1 and Y_2 , each with three levels $\{1, 2, 3\}$. Right: corresponding potential for $\mathcal{M}^{Y_2=3}$.

our MRF. In this paper, we marginalize it exactly using the variable elimination (VE) algorithm (see e.g. Darwiche, 2009b). This method is differentiable w.r.t. all parameters θ since VE only relies on variable summation and factor multiplication. Any other differentiable inference algorithm could be used as well (for instance, the junction tree algorithm as in Kjaerulff and van der Gaag, 2000). This step, evaluating the function g , is known as the *forward pass* in the neural network literature. Next, we backpropagate the previous operation (a step known as the *backward pass*) to build the gradient ∇g . Crucially, note that backpropagation yields $\partial g / \partial \theta$ for every parameter $\theta \in \theta$ of the network at once, not just an individual θ_i . Last, we obtain parameters c_0, c_i, d_0, d_i as detailed before, and use them to compute the metrics of Sections 2.2.1-2.2.3 for each θ_i .

Note the advantages of this approach as compared to other alternatives. For example, symbolically deriving the gradient of g would be cumbersome and depend on the target network topology and definition of the probability of interest (Darwiche, 2003). Automatic differentiation avoids this by evaluating the gradient numerically using the chain rule. Furthermore, finding the gradient using finite differences would require evaluating g twice per parameter θ_i . In contrast, automatic differentiation only requires a forward and backward pass to find the entire gradient –in our experiments, roughly the time of just two marginalization operations (see below).

3.1.4. Additional one-way information

Although YODO is specifically designed to compute the coefficients of the one-way sensitivity function of Equation (3), it further provides all the information to answer additional sensitivity questions:

- It provides the admissible regions for every parameter in θ concerning the event of interest $Y_O = y_O | \mathbf{Y}_E = \mathbf{y}_E$, since they formally only depend on the coefficients c_0, c_i, d_0, d_i as shown in Equations (18) and (19).

- It can quickly find the subset of parameters that do affect the output probability of interest. This is usually called the *parameter sensitivity set* (Coupé and van der Gaag, 2002). This consists of the parameters θ_i for which c_i and/or d_i are non-zero.
- It identifies whether a parameter change leads to a monotonically increasing or decreasing sensitivity function, as already addressed in Bolt and Renooij (2017). Again this can be straightforwardly derived by checking the sign of $c_i d_0 - c_0 d_i$: see Equation (6).

3.2. YODO for multi-way sensitivity analysis

Although there would be no difficulty in conceptually considering simultaneous variations of multiple parameters, we restrict our attention to 2-way sensitivity analyses where only pairs of parameters are varied. This is because: (i) sensitivity functions cannot be visualized in higher dimensions; (ii) the number of groups of parameters grows exponentially; (iii) most critically, the associated measures are challenging to interpret, similar to higher-order interactions in standard statistical models (see e.g. Hayes et al., 2012).

The 2-way version of the sensitivity function considered before would entail computing the unknowns c_{12} and d_{12} from Equation 12. This can be achieved by computing the Hessian, rather than the gradient, in the previous calculations. This can be accomplished via backpropagation too Rojas (1996) and is supported in most modern autodifferentiation packages. However, the sheer size of the Hessian (up to 10^{10} in the networks considered in Sec. 4.1) would make the interpretation of such indices a challenge of its own.

Therefore, we advocate that the maximum n -way sensitivity value is the most valuable and versatile tool for multi-way sensitivity analysis. From its definition in Equation (13), it is clear that it can be instantaneously computed for a specific combination of parameters θ_n once the YODO algorithm has been run. Still, even when focusing on $n = 2$, the possible θ_n can become overwhelmingly large for medium-sized BNs. To address this, we introduce an algorithm to obtain the top K sv_{\max} pairs efficiently by noting that parameters θ contribute to Equation (14) independently from each other. We use a priority queue and proceed in a dynamic programming fashion, whereby we start with a pool \mathcal{P} of K best candidates and keep track of $\max_j sv_{\max}^{\theta_i, \theta_j}$ for all $i \in \mathcal{P}$. The top K pairs are guaranteed to be found after K steps. The algorithm relies on sorting n elements and on K insertions and deletions on the queue and runs in $O(n \log n + K^2 \log K)$ operations. See Algorithm 1 for all details.

Algorithm 1: Algorithm to find the top K maximum 2-way sensitivity values $sv_{\max}^{\theta_i, \theta_j}$ for any $i, j \in [n]$.

```

1: // Gather contributions to Eq. 14 from every BN parameter  $\theta_i$ 
2:  $v \leftarrow$  empty vector
3: for  $i \leftarrow 1$  to  $n$  do
4:    $v_i \leftarrow (c_i d_0^i - c_0^i d_i)^2$ 
5: end for
6:  $v \leftarrow$  sortDescending( $v$ )
7:
8: // Populate the queue with  $K$  initial candidates
9:  $q \leftarrow$  empty priority queue
10: for  $i \leftarrow 1$  to  $K$  do
11:    $q.put(\frac{1}{P(\mathbf{Y}_E = \mathbf{y}_E)^2} \sqrt{v_i + v_{i+1}}, i, i + 1)$  // First element acts as queue's key
12: end for
13:
14: // Read the queue's largest  $K$  keys while updating it
15:  $w \leftarrow$  empty vector
16: for  $k \leftarrow 1$  to  $K$  do
17:    $(v, i, j) \leftarrow q.get()$ 
18:    $w_k \leftarrow v$ 
19:   if  $j < n$  then
20:     // Insert next pair candidate
21:      $q.put(\frac{1}{P(\mathbf{Y}_E = \mathbf{y}_E)^2} \sqrt{v_i + v_{j+1}}, i, j + 1)$ 
22:   end if
23: end for
24: return  $w$ 

```

3.3. Implementation

In order to perform variable elimination efficiently, we note that the problem of graphical model marginalization is equivalent to that of tensor network contraction (Robeva and Seigal, 2018), and use the library *opt_einsum* (Smith and Gray, 2018) which offers optimized heuristics for the latter. As backend, we use the state-of-the-art machine learning library *PyTorch* (Paszke et al., 2019), version 1.13.1, to do all operations between tensors and then perform backpropagation on them. We use *pgmpy* (Ankan and Panda, 2015) for reading and moralizing BNs.

4. Results

We first study the method’s scalability by testing it on large networks with hundreds of nodes and arcs and up to 10^5 parameters; we then overview the insights revealed by our method when applied to two Bayesian networks. All experiments were run on a 4-core i5-6600 3.3GHz Intel workstation with 16GB RAM.

4.1. Simulation study

First, we run our method over the 10 Bayesian networks considered in Scutari et al. (2019). As a baseline, we use the numerical estimation of each sensitivity value via finite differences, whereby we slightly perturb each parameter θ_i and measure the impact on f . As a probability of interest, we set $P(A = a|B = b)$, where A, B, a, b were two variables, and two levels picked randomly, respectively. Each timing is the average of three independent runs. Results are reported in Table 2, which shows that YODO outperforms the baseline by several orders of magnitude and that computing the most relevant 2-way sensitivity values takes in the order of 2s at most.

4.2. Risk assessment for humanitarian crises and disasters

We next extend the analysis of Ballester-Ripoll and Leonelli (2022b), which only focused on one-way indices, to assess the country-level risk associated with humanitarian crises and disasters. The data was collected from INFORM (INFORM, 2022) and consists of 20 drivers of disaster risk covering natural, human, socio-economic, institutional, and infrastructure factors that influence the country-level risk of a disaster, together with a final country risk index which summarizes how exposed a country is to the possibility of a humanitarian disaster. Table 3 reports an overview of the twenty drivers considered, which cover three main risk dimensions: Hazard and exposure (natural/human); Vulnerability (Socio-economic/Vulnerable groups); Lack of coping capacity (institutional/infrastructure). All variables take values between zero and ten. Using the equal-length method, they have been discretized into three categories (low/0, medium/1, high/2). The dataset comprises 190 countries.

Network	#nodes	#arcs	#parameters	Tree-width	Time (fin. diff.)	Time (junction)	Time (ours)	Time (svmax)
child	20	30	344	3	5.9014	0.1274	0.0261	0.0112
water	32	123	13484	10	246.1836	-	0.0579	0.2426
alarm	37	65	752	4	12.0211	8.4439	0.0393	0.0200
haillfinder	56	99	3741	4	58.2177	-	0.0625	0.0920
hepar2	70	158	2139	6	100.0477	-	0.0891	0.0529
win95pts	76	225	1148	8	38.5732	-	0.0923	0.0319
pathfinder	109	208	97851	6	9254.5002	-	0.1826	1.8486
munin1	186	354	19226	11	148307.4553	-	16.0854	1.1949
andes	223	626	2314	17	238.6340	-	0.3115	0.0706
pigs	441	806	8427	10	1544.1509	-	0.5683	0.2131

Table 2: Our method was applied to 10 Bayesian networks, here sorted by the number of nodes. All times are in seconds. The times for the baseline (third-to-last column) were estimated as the total number of parameters in the network and the time needed to estimate one sensitivity value numerically. Treewidths were found with the *NetworkX* graph library Hagberg et al. (2008). The last column reports the time needed to find the top 20 sv_{\max} pairs based on existing YODO gradients.

Variable	Abbreviation	Risk Dimension	Category
Earthquake	EARTHQUAKE	Hazard and Exposure	Natural
Tsunami	TSUNAMI	Hazard and Exposure	Natural
Flood	FLOOD	Hazard and Exposure	Natural
Tropical Cyclone	TROP_CYC	Hazard and Exposure	Natural
Drought	DROUGHT	Hazard and Exposure	Natural
Epidemic	EPIDEMIC	Hazard and Exposure	Natural
Projected Conflict Risk	PCR	Hazard and Exposure	Human
Current Highly Violent Conflict Intensity	CHVCI	Hazard and Exposure	Human
Development and Deprivation	D.AND.D	Vulnerability	Socio-Economic
Economic Dependency	ECON_DEP	Vulnerability	Socio-Economic
Unprotected People	UNP_PEOPLE	Vulnerability	Vulnerable Groups
Other Vulnerable Groups	OTHER_VULN_GROUPS	Vulnerability	Vulnerable Groups
Children U5	CHILDREN_U5	Vulnerability	Vulnerable Groups
Food Security	FOOD_SEC	Vulnerability	Vulnerable Groups
Recent Shocks	RECENT_SHOCKS	Vulnerability	Vulnerable Groups
Health Conditions	HEALTH_COND	Vulnerability	Vulnerable Groups
Governance	GOVERNANCE	Lack of Coping Capacity	Institutional
Communication	COMMUNICATION	Lack of Coping Capacity	Infrastructure
Physical Infrastructure	PHYS_INFRA	Lack of Coping Capacity	Infrastructure
Access to Health System	ACCESS_TO_HEALTH	Lack of Coping Capacity	Infrastructure

Table 3: Variables considered for the humanitarian network from the INFORM (2022) dataset.

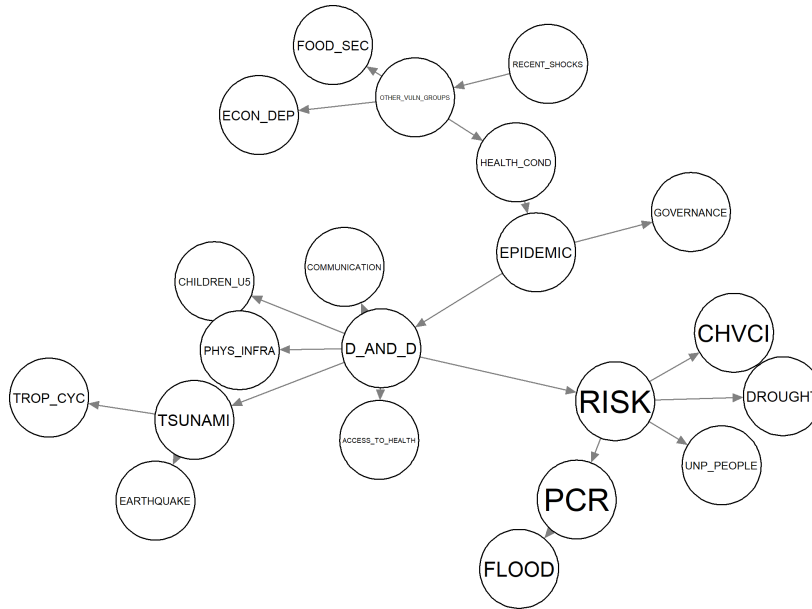


Figure 1: BN learned over the INFORM (2022) dataset for country-level disaster risk.

Similar to Qazi and Simsekler (2021), a BN is learned using the `hc` function of the `bnlearn` package and is reported in Figure 1. A complete interpretation of the learned DAG is beyond the scope of this paper. However, it can be noticed that most risk factors are independent of the overall country-risk given the development and deprivation index (D AND D).

As an illustration of the YODO method, we compute here all sensitivity measures for the conditional probability of a high risk of disaster ($RISK = 2$) conditional on a high risk of flooding ($FLOOD = 2$). Computing all metrics for all 183 network parameters with our method took only 0.055 seconds. The results are reported in Table 4 for the 20 most influential parameters according to the sensitivity value. It can be noticed that the most influential parameters come from the conditional distributions of the overall risk given the development and deprivation index (D AND D), as well as from the conditional distribution of the flooding index given a projected conflict risk index (PCR) equal to low.

As an additional illustration, Figure 2 reports the sensitivity value of the parameters for the output conditional probability of an overall high risk given a high earthquake risk. Blue is associated with positive sensitivity values, and red with negative ones. Out of 183 network parameters, 30 have a sensitivity value of zero, meaning that they do not affect the probability of interest. It can be noticed that the

Parameter	Value	Sensitivity value ↓	Proximity	2 nd deriv.	Largest 1 st deriv.
RISK = high D.AND.D = low	0.0012	0.914	0.056	1.437	0.916
FLOOD = high PCR = low	0.107	0.722	0.0534	4.059	1.475
FLOOD = medium PCR = low	0.469	0.645	0.718	3.238	∞
FLOOD = low PCR = low	0.425	0.645	0.718	3.238	∞
RISK = high D.AND.D = high	0.34	0.555	0.731	0.387	0.714
RISK = high D.AND.D = medium	0.0686	0.467	1.002	0.295	0.488
EPIDEMIC = high HEALTH.COND = low	0.148	0.295	1.167	0.231	0.332
D.AND.D = high EPIDEMIC = medium	0.0742	0.238	1.834	0.133	0.249
PCR = high RISK = medium	0.278	0.226	0.6	0.395	0.394
PCR = high RISK = low	0.0266	0.204	0.694	0.322	0.213
FLOOD = high PCR = high	0.509	0.196	0.475	0.46	1.211
FLOOD = high PCR = medium	0.136	0.167	0.787	0.25	0.206
D.AND.D = high EPIDEMIC = high	0.787	0.159	4.159	0.0459	0.202
D.AND.D = high EPIDEMIC = low	0.0411	0.153	2.984	0.0625	0.156
RISK = low D.AND.D = high	0.0208	0.151	2.319	0.0796	0.274
HEALTH.COND = medium OTHER.VULN.GROUPS = low	0.05	0.151	3.026	0.061	0.154
HEALTH.COND = low OTHER.VULN.GROUPS = low	0.949	0.15	3.036	0.0606	0.154
PCR = low RISK = high	0.00521	0.15	3.023	0.0609	0.236
D.AND.D = medium EPIDEMIC = high	0.176	0.148	5.393	0.0338	0.18
PCR = high RISK = high	0.943	0.148	3.092	0.0588	0.224

Table 4: Four sensitivity metrics for the top 20 parameters of the humanitarian crisis network, when the probability of interest is $P(\text{RISK} = \text{high} | \text{FLOOD} = \text{high})$.

most influential parameters have a positive relationship with the output probability, and almost all are associated with the development and deprivation index.

We further investigate in a 2-way sensitivity analysis the effect of parameters' variations over the same probability $P(\text{RISK} = \text{high} | \text{EARTHQUAKE} = \text{high})$. The 15 largest maximum 2-way sensitivity values are reported in Figure 3. Since these are vector norms, they are always positive irrespective of the relationship between the parameters and the probability of interest. Thus, the coloring should not be interpreted as in Figure 2. Again all parameters associated with the development and deprivation index are the ones that have the most substantial effect on the probability of a country having a high overall risk. Thanks to the efficiency of YODO, these indices are almost instantaneously computed with a total computation time of just 0.047s.

4.3. The role of technology during COVID-19 isolation

The second BN investigates the role of digital communication technology in facilitating the maintenance of meaningful social relationships and promoting the perception of social support during the COVID-19 lockdown. As reported by Gabbiadini et al. (2020), the data was collected through an online questionnaire in March 2020 in Italy, about two weeks from the beginning of the lockdown that the Italian Government adopted for the urgent containment and management of the COVID-19 epidemiological emergency. The data can be downloaded from Gabbiadini (2020) and includes demographic information about 464 individuals, their use of digital communication technologies, and various psychological measures characterizing their emotional status. Each variable is discretized into either two, three, or four levels




Figure 2: Top 20 most influential parameters for the humanitarian crisis network, color-coded by the sign of $f'(\theta_i)$. The probability of interest is $P(\text{RISK} = \text{high} \mid \text{EARTHQUAKE} = \text{high})$. Total computation time: 0.038s.




Figure 3: Top 15 most influential pairs of parameters for the humanitarian crisis network according to sv_{\max} . The probability of interest is $P(\text{RISK} = \text{high} \mid \text{EARTHQUAKE} = \text{high})$. Total computation time: 0.033s.

Variable	Meaning	Group	Levels
AGE	age of respondent	demographic	< 25/≥ 25 (0/1)
GENDER	gender of respondent	demographic	male/female(0/1)
REGION	region of residence	demographic	Lombardy/other(0/1)
OUTSIDE	times outside per week	demographic	0/1/≥2(0/1/2)
SQUARE_METERS	home square meters	demographic	<80/≥80(0/1)
FAMILY_SIZE	number of individuals at home	demographic	1/2/≥ 3(0/1/2)
DAYS_ISOLATION	days since lockdown	demographic	0-10/11-20/>20(0/1/2)
OCCUPATION	occupation	demographic	Other/Smartworking/Student/ Office Work(0/1/2/3)
TECH_FUN_PQ	use of communication technology for fun pre-quarantine	technology	low/medium/high(0/1/2)
TECH_FUN_Q	use of communication technology for fun in quarantine	technology	low/medium/high(0/1/2)
TECH_WORK_PQ	use of communication technology for work pre-quarantine	technology	low/high(0/1)
TECH_WORK_Q	use of communication technology for work in quarantine	technology	low/high(0/1)
ANXIETY	level of anxiety	psychology	low/medium/high(0/1/2)
ANG_IRR	perceived level of anger/irritability	psychology	low/medium/high(0/1/2)
BELONGINGNESS	how often the word we is used	psychology	low/medium/high(0/1/2)
BOREDOM	level of boredom	psychology	low/medium/high(0/1/2)
LONELINESS	perceived loneliness	psychology	low/medium/high(0/1/2)
SOCIAL	perceived social support	psychology	low/medium/high(0/1/2)

Table 5: Variables considered for the COVID-19 network from Gabbiadini et al. (2020).

using either the equal frequency method or some ad-hoc thresholds to optimize the meaning of the levels. Details are reported in Table 5.

A BN is learned for this dataset using 1000 bootstrap repetitions of a tabu search algorithm and keeping the edges that have appeared more than 50% of the times. Furthermore, edges from the psychological measures to the technological and demographic variables were forbidden. Similarly, no edges from the technological to the demographic variables were allowed. These choices were motivated by learning a network whose connections could have a more natural, causal interpretation. Figure 4 reports the learned BN. The two variables connected with psychological measures are age and gender. In particular, given the age of an individual, all other demographic characteristics (except gender) are irrelevant to predict his psychological status. The network, therefore, seems to suggest that age was the main driver for the psychological status of individuals in lockdown.

In this second application, we showcase the computation of the admissible regions using the YODO algorithm. Given a low level of the loneliness index, individuals were most likely spending much time interacting remotely for work during the quarantine (TECH_WORK_Q). Table 6 reports the limits of the admissible regions and other measures ordered from the narrowest interval. The admissible region does not have width one in six cases, all coming from the PMF of TECH_WORK_Q or OCCUPATION. This suggests that the data strongly supports the hypothesis that

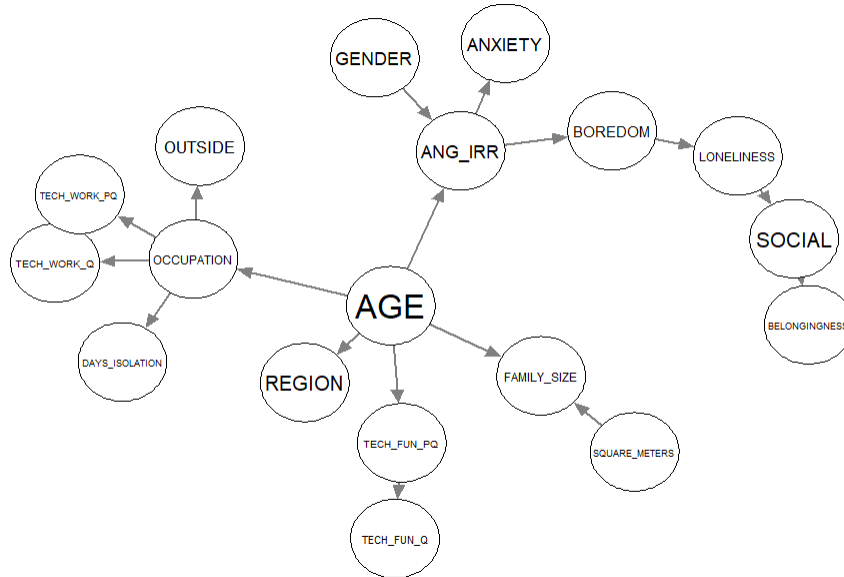


Figure 4: BN learned over the COVID-19 dataset from Gabbiadini et al. (2020).

individuals who did not feel lonely had many online work connections during the lockdown.

As a second illustration, we consider an individual’s age, given that he felt very lonely during the lockdown. The BN suggests the most likely value was of individuals older than 24 years old. Table 7 shows the admissible regions for the network parameters and shows that the network is way less robust for this hypothesis. Admissible regions are much narrower, having a width equal to 0.04 for two parameters. It can also be noticed that parameters with narrow admissible regions come from many different PMFs. Therefore, minor variations in the network parameters would make individuals with less than 25 years more likely to have high levels of loneliness.

5. Discussion

We demonstrated the use of automatic differentiation in BNs and, more specifically, in studying how sensitive they are to parameter variations. The novel algorithms are freely available in Python and are planned to be included in the next release of the `bnmonitor` R package (Leonelli et al., 2021). Their efficiency was demonstrated through a simulation study. Two critical applications in humanitarian crises and studying the psychological effects of isolation during the COVID-19

Parameter	Value	Sens. value	Proximity	AR (lower)	AR (upper)
TECH_WORK_Q = 1 TECH_WORK_PQ = 1, OCCUPATION = 1	0.96	0.21	$4.54 \cdot 10^6$	0.56	1.0
OCCUPATION = 0 AGE = 1	0.37	0.28	$9.91 \cdot 10^6$	0	0.67
OCCUPATION = 1 AGE = 1	0.4	0.28	$9.78 \cdot 10^6$	0.1	1.0
TECH_WORK_Q = 1 TECH_WORK_PQ = 0, OCCUPATION = 2	0.41	0.27	$2.02 \cdot 10^7$	0.093	1.0
TECH_WORK_Q = 0 TECH_WORK_PQ = 0, OCCUPATION = 2	0.59	0.16	N/A	0.068	1.0
TECH_WORK_Q = 1 TECH_WORK_PQ = 1, OCCUPATION = 2	0.75	0.12	$5.49 \cdot 10^7$	0.041	1.0
AGE = 0	0.45	0.11	4.36	0	1.0
TECH_FUN_PQ = 1 AGE = 1	0.26	$4.0 \cdot 10^{-9}$	$2.1 \cdot 10^7$	0	1.0
TECH_FUN_PQ = 2 AGE = 0	0.36	$2.39 \cdot 10^{-8}$	N/A	0	1.0
TECH_FUN_PQ = 2 AGE = 1	0.51	$8.01 \cdot 10^{-9}$	$1.05 \cdot 10^7$	0	1.0
TECH_FUN_Q = 0 TECH_FUN_PQ = 0	0.47	$2.18 \cdot 10^{-8}$	$4.19 \cdot 10^7$	0	1.0
TECH_FUN_Q = 0 TECH_FUN_PQ = 1	0.27	$1.39 \cdot 10^{-8}$	$4.19 \cdot 10^7$	0	1.0
TECH_FUN_Q = 0 TECH_FUN_PQ = 2	0.12	0	N/A	0	1.0
TECH_FUN_Q = 1 TECH_FUN_PQ = 0	0.34	$7.92 \cdot 10^{-9}$	$2.1 \cdot 10^7$	0	1.0
TECH_FUN_Q = 1 TECH_FUN_PQ = 1	0.41	$1.39 \cdot 10^{-8}$	$4.19 \cdot 10^7$	0	1.0

Table 6: COVID network for the probability of interest $P(\text{TECH_WORK_Q} = \text{high} \mid \text{LONELINESS} = \text{low})$: sensitivity metrics for the 15 parameters with the smallest admissible region. Total computation time: 0.071s.

Parameter	Value	Sens. value	Proximity	AR (lower)	AR (upper)
DAYS_ISOLATION = 0 OCCUPATION = 3	1.0	0.028	11.13	0.96	1.0
OUTSIDE = 2 OCCUPATION = 3	1.0	0.028	11.13	0.96	1.0
ANG_IRR = 2 AGE = 0, GENDER = 0	0.27	0.047	8.4	0	0.29
BOREDOM = 0 ANG_IRR = 1	0.24	0.019	3.4	0	0.3
ANG_IRR = 0 AGE = 1, GENDER = 1	0.3	0.13	2.43	0	0.31
LONELINESS = 1 BOREDOME = 0	0.33	0.014	7.4	0	0.41
LONELINESS = 0 BOREDOME = 1	0.32	0.011	4.28	0	0.42
AGE = 1	0.55	1.01	0.023	0.54	1.0
AGE = 0	0.45	0.59	1.3	0	0.46
GENDER = 1	0.75	0.0048	4.6	0.52	1.0
GENDER = 0	0.25	0.0048	4.6	0	0.48
ANG_IRR = 1 AGE = 1, GENDER = 1	0.42	0.018	23.86	0	0.49
ANG_IRR = 2 AGE = 0, GENDER = 1	0.49	0.13	2.42	0	0.5
LONELINESS = 1 BOREDOME = 1	0.47	0.011	4.28	0	0.57
BOREDOME = 1 ANG_IRR = 1	0.5	0.016	4.12	0	0.58

Table 7: COVID network for the probability of interest $P(\text{AGE} \geq 25 \mid \text{LONELINESS} = \text{high})$: sensitivity metrics for the 15 parameters with the smallest admissible region. Total computation time: 0.11s.

pandemic illustrate their use in practice.

Although YODO is specifically designed to compute the coefficients of the one-way sensitivity function in Equation 3, we demonstrated in this paper how it could be used to answer a variety of sensitivity queries, for instance, admissible regions and the identification of the parameter sensitivity set. Importantly, YODO also provides the basis for multi-way sensitivity analyses, and we demonstrated their feasibility in practice.

Future Work

The YODO algorithm introduced here is designed explicitly for BN models, but it could also be adapted to work with other graphical models. The study of context-specific independence has been shown to increase the efficiency of various inferential tasks often, and thus we may expect that it could also speed up YODO. Therefore, we plan to adapt it to work over graphical models embedding non-symmetric types of independence, as, for instance, staged trees (Carli et al., 2022; Smith and Anderson, 2008), whose sensitivity functions have also been studied (Leonelli, 2019). Another avenue of research is the adaptation of YODO to work for sum-product networks (Poon and Domingos, 2011; Sánchez-Cauce et al., 2022), a different representation of a factorization of a joint probability distribution, which has become increasingly popular in the past few years.

Although YODO makes various types of multi-way sensitivity analysis feasible, they are still a local approach to investigate the combined effect of parameters' variations on probabilities of interest. Recently, it has been shown that the computation of Sobol indices, a global sensitivity index, is feasible in sensitivity to evidence analyses (Ballester-Ripoll and Leonelli, 2022a). We are currently investigating algorithms to globally assess the effect of the various parameters of a BN and consequently compute their associated Sobol indices.

References

- Ankan, A., Panda, A., 2015. pgmpy: Probabilistic graphical models using python, in: Proceedings of the 14th Python in Science Conference (SCIPY 2015), Citeseer.
- Ballester-Ripoll, R., Leonelli, M., 2022a. Computing Sobol indices in probabilistic graphical models. *Reliability Engineering & System Safety* 225, 108573.
- Ballester-Ripoll, R., Leonelli, M., 2022b. You only derive once (YODO): Automatic differentiation for efficient sensitivity analysis in Bayesian networks, in: International Conference on Probabilistic Graphical Models, PMLR. pp. 169–180.

- Bielza, C., Larranaga, P., 2014. Discrete Bayesian network classifiers: A survey. *ACM Computing Surveys* 47, 1–43.
- Bolt, J.H., van der Gaag, L.C., 2017. Balanced sensitivity functions for tuning multi-dimensional Bayesian network classifiers. *International Journal of Approximate Reasoning* 80, 361–376.
- Bolt, J.H., Renooij, S., 2014. Local sensitivity of Bayesian networks to multiple simultaneous parameter shifts, in: *European Workshop on Probabilistic Graphical Models*, Springer. pp. 65–80.
- Bolt, J.H., Renooij, S., 2017. Structure-based categorisation of Bayesian network parameters, in: *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Springer. pp. 83–92.
- Carli, F., Leonelli, M., Riccomagno, E., Varando, G., 2022. The R package stagedtrees for structural learning of stratified staged trees. *Journal of Statistical Software* 102, 1–30.
- Castillo, E., Gutiérrez, J.M., Hadi, A.S., 1997. Sensitivity analysis in discrete Bayesian networks. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 27, 412–423.
- Chan, H., Darwiche, A., 2004. Sensitivity analysis in Bayesian networks: From single to multiple parameters, in: *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pp. 67–75.
- Chan, H., Darwiche, A., 2005. A distance measure for bounding probabilistic belief change. *International Journal of Approximate Reasoning* 38, 149–174.
- Coupé, V.M., van der Gaag, L.C., 2002. Properties of sensitivity analysis of Bayesian belief networks. *Annals of Mathematics and Artificial Intelligence* 36, 323–356.
- Coupé, V.M., van der Gaag, L.C., Habbema, J.D.F., 2000. Sensitivity analysis: An aid for belief-network quantification. *The Knowledge Engineering Review* 15, 215–232.
- Darwiche, A., 2003. A differential approach to inference in Bayesian networks. *Journal of the ACM* 50, 280–305.
- Darwiche, A., 2009a. *Modeling and reasoning with Bayesian networks*. Cambridge University Press.

- Darwiche, A., 2009b. Modeling and reasoning with Bayesian networks. Cambridge University Press.
- Douglas-Smith, D., Iwanaga, T., Croke, B.F., Jakeman, A.J., 2020. Certain trends in uncertainty and sensitivity analysis: An overview of software tools and techniques. *Environmental Modelling & Software* 124, 104588.
- Friedman, N., Geiger, D., Goldszmidt, M., 1997. Bayesian network classifiers. *Machine learning* 29, 131–163.
- van der Gaag, L., Renooij, S., 2001. Analysing sensitivity data from probabilistic networks, in: *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, pp. 530–537.
- van der Gaag, L.C., Renooij, S., Coupé, V.M., 2007. Sensitivity analysis of probabilistic networks, in: *Advances in Probabilistic Graphical Models*. Springer, pp. 103–124.
- Gabbiadini, A., 2020. The mitigating role of digital communication technologies on negative affect during the covid-19 outbreak in italy. URL: [doi:10.20366/unimib/unidata/SN223-1.0](https://doi.org/10.20366/unimib/unidata/SN223-1.0). uniData - Bicocca Data Archive, Milan. Study Number SN223. Data file version 1.0.
- Gabbiadini, A., Baldissarri, C., Durante, F., Valtorta, R.R., De Rosa, M., Gallucci, M., 2020. Together apart: The mitigating role of digital communication technologies on negative affect during the covid-19 outbreak in italy. *Frontiers in Psychology* 11, 554678.
- Goerlandt, F., Islam, S., 2021. A Bayesian Network risk model for estimating coastal maritime transportation delays following an earthquake in British Columbia. *Reliability Engineering & System Safety* 214, 107708.
- Hagberg, A.A., Schult, D.A., Swart, P.J., 2008. Exploring network structure, dynamics, and function using NetworkX, in: *Proceedings of the 7th Python in Science Conference*, pp. 11–15.
- Hayes, A.F., Glynn, C.J., Huges, M.E., 2012. Cautions regarding the interpretation of regression coefficients and hypothesis tests in linear models with interactions. *Communication Methods and Measures* 6, 1–11.
- INFORM, 2022. Index for risk management. Retrieved from <https://drmkc.jrc.ec.europa.eu/inform-index>.

- Kjaerulff, U., van der Gaag, L., 2000. Making sensitivity analysis computationally efficient, in: Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence, pp. 317–325.
- Kwisthout, J., van der Gaag, L., 2008. The computational complexity of sensitivity analysis and parameter tuning, in: Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence, pp. 349–356.
- Laskey, K.B., 1995. Sensitivity analysis for probability assessments in Bayesian networks. *IEEE Transactions on Systems, Man, and Cybernetics* 25, 901–909.
- Leonelli, M., 2019. Sensitivity analysis beyond linearity. *International Journal of Approximate Reasoning* 113, 106–118.
- Leonelli, M., Görgen, C., Smith, J.Q., 2017. Sensitivity analysis in multilinear probabilistic models. *Information Sciences* 411, 84–97.
- Leonelli, M., Ramanathan, R., Wilkerson, R.L., 2021. Sensitivity and robustness analysis in Bayesian networks with the bnmonitor R package. [arXiv:2107.11785](https://arxiv.org/abs/2107.11785).
- Leonelli, M., Riccomagno, E., 2022. A geometric characterisation of sensitivity analysis in monomial models. *International Journal of Approximate Reasoning* 151, 64–84.
- Li, C., Mahadevan, S., 2018. Sensitivity analysis of a Bayesian network. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems Part B: Mechanical Engineering* 4.
- Makaba, T., Doorsamy, W., Paul, B.S., 2021. Bayesian network-based framework for cost-implication assessment of road traffic collisions. *International Journal of Intelligent Transportation Systems Research* 19, 240–253.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An imperative style, high-performance deep learning library, in: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc., pp. 8024–8035.
- Pearl, J., 1988. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufmann.

- Poon, H., Domingos, P., 2011. Sum-product networks: A new deep architecture, in: 2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops), IEEE. pp. 689–690.
- Qazi, A., Simsekler, M.C.E., 2021. Assessment of humanitarian crises and disaster risk exposure using data-driven Bayesian networks. *International Journal of Disaster Risk Reduction* 52, 101938.
- Renooij, S., 2014. Co-variation for sensitivity analysis in Bayesian networks: Properties, consequences and alternatives. *International Journal of Approximate Reasoning* 55, 1022–1042.
- Robeva, E., Seigal, A., 2018. Duality of graphical models and tensor networks. *Information and Inference: A Journal of the IMA* 8, 273–288.
- Rohmer, J., 2020. Uncertainties in conditional probability tables of discrete Bayesian belief networks: A comprehensive review. *Engineering Applications of Artificial Intelligence* 88, 103384.
- Rojas, R., 1996. *Neural Networks - A Systematic Introduction*. Springer-Verlag, Berlin. URL: http://www.inf.fu-berlin.de/inst/ag-ki/rojas_home/pmwiki/pmwiki.php?n=Books.NeuralNetworksBook.
- Sánchez-Cauce, R., París, I., Díez, F.J., 2022. Sum-product networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 3821–3839.
- Scutari, M., Graafland, C.E., Gutiérrez, J.M., 2019. Who learns better Bayesian network structures: Accuracy and speed of structure learning algorithms. *International Journal of Approximate Reasoning* 115, 235–253.
- Smith, D.G.A., Gray, J., 2018. opt_einsum - a Python package for optimizing contraction order for einsum-like expressions. *Journal of Open Source Software* 3, 753.
- Smith, J.Q., Anderson, P.E., 2008. Conditional independence and chain event graphs. *Artificial Intelligence* 172, 42–68.
- Zio, E., Mustafayeva, M., Montanaro, A., 2022. A Bayesian belief network model for the risk assessment and management of premature screen-out during hydraulic fracturing. *Reliability Engineering & System Safety* 218, 108094.