

Subjective Networks: Prospectives and Challenges *

Magdalena Ivanovska¹, Audun Jøsang¹, Lance Kaplan², and Francesco Sambo³

¹University of Oslo, Norway

²US Army Research Lab, Adelphi, MD, US

³University of Padova, Italy

Abstract

Subjective logic is a formalism for representing and reasoning under uncertain probabilistic information, with an explicit treatment of the uncertainty about the probability distributions. We introduce *subjective networks* as graph-based structures that generalize Bayesian networks to the theory of subjective logic. We discuss the prospectives of the subjective networks representation and the challenges of reasoning with them.

1 Introduction

Subjective logic [Jøsang, 2001] is a formalism for representing and reasoning under uncertain probabilistic information. The basic entities in subjective logic are *subjective opinions* on random variables. A subjective opinion includes a *belief mass distribution* over the states of the variable complemented with an *uncertainty mass*, reflecting a current analysis of the probability distribution of the variable by an expert, based on a test, etc; and a *base rate* probability distribution of the variable, reflecting a domain knowledge that is relevant to the current analysis. A subjective opinion can always be *projected* onto a single probability distribution, but this necessarily removes information about the uncertainty mass.

While a probability distribution itself represents uncertainty about the value of the variable, a subjective opinion represents a second-order uncertainty, i.e. uncertainty about the probability distribution. The latter is further formalized by establishing a correspondence between subjective opinions and Dirichlet probability density functions [Jøsang and McAnally, 2004].

Conditional reasoning with subjective opinions has been explored for the case of two variables, resulting in the definition of *deduction* and *abduction* operations for multinomial variables [Jøsang, 2008]. An alternative approach to deduction based on the Dirichlet model is explored in [Kaplan *et al.*, 2013].

This paper attempts to address the conditional reasoning with subjective opinions in general, introducing *subjective*

networks as graph-based structures that generalize Bayesian networks to the theory of subjective logic.

A Bayesian network [Pearl, 1988] is a compact representation of a joint probability distribution of a set of random variables in the form of directed acyclic graph and a set of conditional probability distributions associated with each node. The goal of inference in Bayesian networks is to derive the conditional probability distribution of any set of (target) variables in the network, given that the values of any other set of (evidence) variables have been observed. Bayesian networks reasoning algorithms provide a way to propagate the probabilistic information through the graph, from the evidence to the target. Bayesian networks are a powerful tool for modelling and inference of various situations involving probabilistic information about a set of variables, and thus form a base for developing tools with applications in many areas like medical diagnostics, risk management, etc.

One serious limitation of the Bayesian networks reasoning algorithms is that all the input conditional probabilities must be assigned precise values in order for the inference algorithms to work and the model to be analysed. This is problematic in situations where probabilities can not be reliably elicited and one needs to do inference with uncertain or incomplete probabilistic information, inferring the most accurate conclusions possible. Subjective opinions can represent uncertain probabilistic information of any kind (minor or major imprecision, and even total ignorance), by varying the uncertainty mass between 0 and 1.

A straightforward generalization of Bayesian networks in subjective logic retains the network structure and replaces conditional probability distributions with conditional subjective opinions at every node of the network. We call this a *Bayesian subjective network* and consider the reasoning in it as a generalization of classical Bayesian reasoning, where the goal is to obtain a subjective opinion on the target given the evidence. The evidence in this case can be an instantiation of values, but also a subjective opinion itself. In most of the cases, the inference in subjective Bayesian networks remains a challenge, since subjective opinions do not enjoy all the nice properties of probability distributions and, in particular, the notions of conditioning and joint opinion are not defined in general in subjective logic. We also discuss representation and inference with *fused subjective networks*, where the graph follows the available input information as associated with the

*This research is conducted within the project ABRI partially funded by the US Army Research Laboratory.

arrows rather than the nodes, and where information coming from multiple paths to the same node is combined by *fusion* operation. We give an example of modelling with subjective networks through the special case of the *naïve Bayes subjective network*, which can be considered to belong to both the Bayesian and the fused subjective networks type.

The paper is structured as follows: In Section 2 we first review the necessary preliminaries from probability theory and Bayesian networks. Then we introduce subjective opinions on random variables, and their correspondence of with the multinomial Dirichlet model. Section 3 introduces subjective networks representation. In Section 4 we introduce the types of inference problems that can be distinguished in subjective networks and discuss potential solutions. Section 5 presents an alternative approach to inference in subjective networks that builds upon the Dirichlet representation of subjective opinions. In Section 6 we conclude the paper.

2 Preliminaries

2.1 Bayesian Networks

We assume a simplified definition of *random variable* as a variable that takes its values with certain probabilities. More formally, let X be a variable with a *domain* (set of values, states of the variable) \mathbb{X} . A *probability distribution* p of X is a function $p : \mathbb{X} \rightarrow [0, 1]$, such that:

$$\sum_{x \in \mathbb{X}} p(x) = 1. \quad (1)$$

$p(x)$ is the probability that the variable X takes the value x .

Let $V = \{X_1, \dots, X_n\}$ be the set of all random variables that are of interest in a given context. A *joint probability distribution* of the variables V is a probability distribution defined on the Cartesian product of $\mathbb{X}_1, \dots, \mathbb{X}_n$:

$$\sum_{x_1 \in \mathbb{X}_1} \dots \sum_{x_n \in \mathbb{X}_n} p(x_1, \dots, x_n) = 1. \quad (2)$$

In general we will talk about sets of variables, subsets of V , identifying the variables themselves with the singleton subsets. As standard in Bayesian networks literature, we use the notation of a random variable for a set of variables, making the obvious identifications (see [Pearl, 1988]).

Given a joint probability distribution p of the variables in V , and a set of variables $Y \subseteq V$, the *marginal probability distribution* of Y , $p : \mathbb{Y} \rightarrow [0, 1]$ is given by:

$$p(y) = \sum_{x \in \mathbb{X}, X=V \setminus Y} p(y, x). \quad (3)$$

Given two sets of variables X and Y , a *conditional probability distribution* of Y given that X takes the value x , $p(Y|x)$, is a function from \mathbb{Y} to $[0, 1]$ defined by:

$$p(y|x) = \frac{p(y, x)}{p(x)}. \quad (4)$$

$p(y|x)$ is the conditional probability that Y takes the value y , given that the value of X is x .

A set of variables X is conditionally independent of a set of variables Y given the set of variables Z , written $I(X, Y|Z)$ if the following holds:

$$p(x|y, z) = p(x|z) \text{ whenever } p(y, z) > 0, \quad (5)$$

for every choice of assignments x, y , and z .

A Bayesian network [Pearl, 1988] with n variables is a directed acyclic graph (DAG) with random variables X_1, \dots, X_n as nodes, and a set of conditional probability distributions $p(X_i|Pa(X_i))$, associated with each node X_i , containing one conditional probability distribution $p(X_i|pa(X_i))$ of X_i for every assignment of values $pa(X_i)$ to its parent nodes. If we assume that the Markov property holds: Every node is conditionally independent on its non-descendants given its parents,

$$I(X_i, ND(X_i)|Pa(X_i)), \quad (6)$$

for the given DAG and the joint distribution p , then p is determined by:

$$p(x_1, \dots, x_n) = \prod_{i=1}^n p(x_i|pa(X_i)), \quad (7)$$

where $pa(X_i)$ is the instantiation of the parents of X_i that corresponds to the tuple (x_1, \dots, x_n) .

The general goal of inference in Bayesian networks is to be able to derive the probability $p(y|x)$, for every choice of values of arbitrary sets of variables X and Y , in an efficient way compatible with the network's topology.

2.2 Subjective Opinions

In this section we review the basic notions related to multinomial and hyper opinions on random variables.

Let X be a random variable. A *multinomial subjective opinion* on X [Jøsang, 2008] is a tuple:

$$\omega_X = (b_X, u_X, a_X), \quad (8)$$

where $b_X : \mathbb{X} \rightarrow [0, 1]$ is a *belief mass distribution*, $u_X \in [0, 1]$ is an *uncertainty mass*, and $a_X : \mathbb{X} \rightarrow [0, 1]$ is a *base rate distribution*, satisfying the following additivity constraints:

$$u_X + \sum_{x \in \mathbb{X}} b_X(x) = 1, \quad (9)$$

$$\sum_{x \in \mathbb{X}} a_X(x) = 1. \quad (10)$$

The beliefs and the uncertainty mass reflect the results of a current analysis of the random variable applying expert knowledge, experiments, or a combination of the two. $b_X(x)$ is the belief that X takes the value x expressed as a degree in $[0, 1]$. It represents the amount of experimental or analytical evidence in favour of x . u_X is a single value, representing the degree of uncertainty about the distribution of X . It represents lack of evidence due to lack of knowledge or expertise, or insufficient experimental analysis. The base rate a_X is simply a probability distribution of X that represents domain knowledge relevant to the current analysis.

For example, a GP wants to determine whether a patient suffers from depression through a series of different tests.

Based on the test results, the GP concludes that the collected evidence is 10% inconclusive, but is still two times more in support of the diagnosis that the patient suffers from depression than of the opposite one. As a result the GP assigns 0.6 belief mass to the diagnosis that the patient suffers from depression and 0.3 belief mass to the opposite diagnosis, complemented by 0.1 uncertainty mass. The probability that a random person in the population suffers from depression is 5% and this fact determines the base rate distribution in the GP's subjective opinion on the condition of the patient.

In some cases of modelling, it is useful to be able to distribute belief mass to subsets of \mathbb{X} as well. This leads to generalization of multinomial subjective opinions to *hyper opinions*, which distribute the belief mass over the reduced power set of \mathbb{X} (*hyperdomain* of X), $\mathcal{R}(\mathbb{X}) = \mathcal{P}(\mathbb{X}) \setminus \{\mathbb{X}, \emptyset\}$:

$$b_X : \mathcal{R}(\mathbb{X}) \rightarrow [0, 1], \quad (11)$$

and u_X is a value from $[0, 1]$, such that the following holds:

$$u_X + \sum_{x \in \mathcal{R}(\mathbb{X})} b_X(x) = 1. \quad (12)$$

a_X is again a probability distribution of X , defined on \mathbb{X} .¹

$b_X(x)$ represents the belief that the value of X is (in the set) $x \in \mathcal{R}(\mathbb{X})$, and represents the amount of evidence that supports exactly x .²

A subjective opinion in which $u_X = 0$, i.e. an opinion without uncertainty mass, is called a *dogmatic opinion*. Dogmatic multinomial opinions correspond to probability distributions. A dogmatic opinion for which $b_X(x) = 1$, for some $x \in \mathbb{X}$, is called an *absolute opinion* and denoted by ω_X^x . Absolute multinomial opinions correspond to instantiating values of variables. In contrast, an opinion for which $u_X = 1$, and consequently $b_X(x) = 0$, for every $x \in \mathcal{R}(\mathbb{X})$, i.e. an opinion with complete uncertainty, is called a *vacuous opinion*. Vacuous opinions correspond to complete ignorance about the probability distribution of the variable.

A multinomial opinion ω_X is "projected" to a probability distribution $P_{\omega_X} : \mathbb{X} \rightarrow [0, 1]$, defined in the following way:

$$P_{\omega_X}(x) = b_X(x) + a_X(x) u_X. \quad (13)$$

We call the function P_{ω_X} a *projected probability distribution* of ω_X . According to Eq.(13), $P_{\omega_X}(x)$ is the belief mass in support of x increased by the portion of the base rate of x that is determined by u_X . In that way, it gives an estimate of the probability of x which varies from the base rate value, in the case of complete ignorance, to the actual probability in the case of zero uncertainty.

For hyper opinions, the definition of projected probability distribution is generalized as follows:

$$P_{\omega_X}(x) = \sum_{x' \in \mathcal{R}(\mathbb{X})} a_X(x|x') b_X(x') + a_X(x) u_X, \quad (14)$$

¹We abuse the notation by using the same type of letters for both elements of \mathbb{X} and elements of $\mathcal{R}(\mathbb{X})$.

²If we think of u_X as of an amount of evidence assigned to the whole domain \mathbb{X} , then b_X and u_X correspond to a *basic belief assignment* [Shafer, 1976]. However, u_X is a measure for lack of evidence, not a belief, as will be further clarified in the next section.

for $x \in \mathbb{X}$, where $a_X(x|x')$ is the conditional probability of x given x' , if a_X is extended to $\mathcal{P}(X)$ additively.³ If we denote the sum in Eq.(14) by b'_X :

$$b'_X(x) = \sum_{x' \in \mathcal{R}(\mathbb{X})} a_X(x|x') b_X(x'), \quad (15)$$

it is easy to check that $b'_X : \mathbb{X} \rightarrow [0, 1]$, together with u_X , satisfies the additivity property in Eq.(12), i.e. $\omega'_X = (b'_X, u_X, a_X)$ is a multinomial opinion. From Eq.(14) and Eq.(15) we obtain $P_{\omega_X} = P_{\omega'_X}$. This means that every hyper opinion can be approximated with a multinomial opinion which has the same projected probability distribution as the initial hyper one.

2.3 Subjective Opinions as Dirichlet pdfs

In this section we describe the correspondence between multinomial opinions and multinomial Dirichlet models.

Let $\mathbf{p} = (p_1, \dots, p_k)$ be the probability distribution of the variable X , where $p_i = p(x_i)$. \mathbf{p} is Dirichlet distributed if its probability density function (pdf) has the following form:

$$f_{\alpha}(\mathbf{p}) = \frac{\Gamma\left(\sum_{i=1}^k \alpha_i\right)}{\prod_{i=1}^k \Gamma(\alpha_i)} \prod_{i=1}^k p_i^{\alpha_i-1}, \quad (16)$$

where Γ is the k -dimensional ($k = |\mathbb{X}|$) gamma function and $\alpha = (\alpha_1, \dots, \alpha_k)$ are the parameters of the distribution. The *mean* distribution is determined by $m(x_i) = \alpha_i / \sum_{i=1}^k \alpha_i$.

The *multinomial Dirichlet model* [Gelman and others, 2004] assumes: i) a Dirichlet prior pdf for \mathbf{p} with parameters $\alpha_i^0 = C a_X(x_i)$, where a_X is the mean distribution and C is a *prior strength* determining the amount of evidence needed to overcome the prior; ii) multinomial sampling ($r(x_i) \mid i = 1, \dots, k$), $N = \sum_{i=1}^k r(x_k)$, i.e. N observations where x_i is observed $r(x_i)$ times. Then the posterior pdf for \mathbf{p} is also a Dirichlet pdf with the following parameters:

$$\alpha_i = r(x_i) + C a_X(x_i), \quad (17)$$

and the following mean distribution:

$$m(x_i) = \frac{r(x_i) + C a_X(x_i)}{N + C}. \quad (18)$$

The posterior Dirichlet pdf for \mathbf{p} uniquely determines a multinomial opinion $\omega_X = (b_X, u_X, a_X)$, where:

$$\begin{cases} b_X(x_i) = \frac{r(x_i)}{N+C} \\ u_X = \frac{C}{N+C} \end{cases} \quad (19)$$

By the transformation in Eq.(19), the projected probability of the obtained ω_X is equal to the mean of the posterior Dirichlet pdf as given in Eq.(18), which corresponds to the fact that it represents an estimate for the actual distribution \mathbf{p} of X .⁴

³For this conditional probability to be always defined, it is enough to assume $a_X(x_i) > 0$, for every $x_i \in \mathbb{X}$. This amounts to assuming that everything we include in the domain has a non-zero probability of occurrence.

⁴A dogmatic opinion ($u_X = 0$) is obtained when the number of observations converges to infinity. In that case the beliefs converge to the actual \mathbf{p} .

Conversely, a given opinion $\omega_X = (b_X, u_X, a_X)$ and a Dirichlet strength C determine a Dirichlet pdf with parameters α as given by Eq.(17), where:

$$r(x_i) = \frac{Cb_X(x_i)}{u_X} \quad (20)$$

This means that expressing a subjective opinion $\omega_X = (b_X, u_X, a_X)$ is equivalent to expressing the subject's knowledge about the actual probability distribution of X by a multinomial Dirichlet pdf. This correspondence gives a way of eliciting the beliefs and uncertainty in subjective opinions from experimental analysis. However, the base rate as well as its strength C have to be chosen in advance.

3 Subjective Networks Representation

A *subjective network* S_n of n random variables is a directed acyclic graph and sets of subjective opinions associated with it. First we introduce the concepts of joint and conditional subjective opinion, then we introduce two different types of subjective networks, Bayesian and fused subjective networks.

3.1 Conditional and Joint Subjective Opinions

A *joint subjective opinion* on variables $X_1, \dots, X_n, n \geq 2$ is the tuple:

$$\omega_{X_1 \dots X_n} = (b_{X_1 \dots X_n}, u_{X_1 \dots X_n}, a_{X_1 \dots X_n}), \quad (21)$$

where $b_{X_1 \dots X_n} : \mathcal{R}(\mathbb{X}_1 \times \dots \times \mathbb{X}_n) \rightarrow [0, 1]$ and $u_{X_1 \dots X_n} \in [0, 1]$ satisfy the condition from Eq.(12) and $a_{X_1 \dots X_n}$ is a joint probability distribution of X_1, \dots, X_n .

A *marginal opinion* on a set of variables Y , subset of $V = \{X_1, \dots, X_n\}$, is a joint opinion on the variables in Y . The relation between a marginal opinion on the variables Y and a joint opinion on the full set of variables V can not be modelled with an analogue of Eq.(3), but rather with what is known as the *product operation* [Jøsang and McAnally, 2004], where a product of two multinomial opinions on independent random variables is defined as a joint hyper opinion on the Cartesian domain. The definition is generalizable to an arbitrary number of variables and to opinions on sets of variables, under the assumption that the input (sets of) opinions are subjective opinions on probabilistically independent variables.

Given two sets of random variables X and Y , a conditional opinion on Y given that X takes the value x is a subjective opinion on Y defined as a tuple:

$$\omega_{Y|x} = (b_{Y|x}, u_{Y|x}, a_{Y|x}), \quad (22)$$

where $b_{Y|x} : \mathcal{R}(\mathbb{Y}) \rightarrow [0, 1]$ and $u_{Y|x} \in [0, 1]$ satisfy the condition in Eq.(12) and $a_{Y|x} : \mathbb{Y} \rightarrow [0, 1]$ is a probability distribution of Y . We use the notation $\omega_{Y|X}$ for a set of conditional opinions on Y , one for every value of X :

$$\omega_{Y|X} = \{\omega_{Y|x} \mid x \in X\}. \quad (23)$$

There is no relation in subjective logic analogous to Eq.(4) that defines conditional opinions through marginal opinions.

3.2 Bayesian Subjective Networks

A *Bayesian subjective network* of n random variables X_1, \dots, X_n is a directed acyclic graph with one node for each variable and a set of conditional subjective opinions $\omega_{X_i|Pa(X_i)}$ associated with each node X_i , consisting of one conditional opinion $\omega_{X_i|pa(X_i)}$ on X_i , for each instantiation $pa(X_i)$ of its parent nodes $Pa(X_i)$.

A Bayesian subjective network is basically a generalization of a classical Bayesian network where instead of probability distributions associated with the nodes, we have subjective opinion about them. Conversely, every Bayesian subjective network projects to a classical one. Namely, every opinion $\omega_{X_i|pa(X_i)} \in \omega_{X_i|Pa(X_i)}$ projects to a probability distribution $P(X_i|pa(X_i))$. The graph of the given subjective network S_n together with the sets of projected distributions $P(X_i|Pa(X_i)), i = 1, \dots, n$, forms a classical Bayesian network, which we denote by $P(S_n)$ and call a *Bayesian network projection* of the network S_n .

The concepts of joint, marginal, and conditional opinions do not enjoy the same relations as their probabilistic counterparts. Consequently, the joint opinion on X_1, \dots, X_n can not be obtained from the given opinions in the network using the Markov condition, i.e. by an analogue of Eq.(7). Nevertheless, the corresponding projected probabilities are related by the equations in Section 2.1 and can be reasoned about within the Bayesian network projection in the classical way.

The Bayesian subjective networks representation also presumes that Markov independences, hence all the conditional independences embedded in the graph structure of the given DAG (d-separations), hold for the uncertainties of the corresponding opinions: If a set of variables X is conditionally independent of a set of variables Y given the set of variables $Z, I(X, Y|Z)$ then:

$$u(X|yz) = u(X|z), \quad (24)$$

for every choice of values y and z . This assumption can be justified by the fact that the uncertainty mass of a subjective opinion is a parameter that refers to the whole probability distribution. In light of the Dirichlet pdf representation of subjective opinions, a subjective network represents in some sense an ensemble of possible Bayesian networks, where the spread of the distributions is related to the uncertainties. For each distribution in the ensemble, $I(X, Y|Z)$ implies $p(X|yz) = p(X|z)$. Therefore, the spread of $p(X|yz)$ is the same as that of $p(X|z)$.

A subjective network S_n is a graphical representation of uncertain information about probability distributions that combines beliefs and uncertainty, as well as probabilistic information about the knowledge domain in the form of base rate distributions. The base rate distributions in the conditional opinions of a subjective network S_n can be set without constraints and are not necessarily connected by the equations in Section 2.1, i.e. the subjective network may or may not represent a joint probability distribution on the knowledge domain. In the two-node case considered in [Jøsang, 2008] for example, it is assumed $a_{Y|x_i} = a_Y$, for every $x_i \in X$, i.e. that only the unconditional base rate distributions are available.




Figure 1: A three node Naïve Bayes subjective network.

Example: Naïve Bayes Subjective Networks

A specific case which is often modelled in data mining and machine learning is that of a set of variables X_1, \dots, X_n , all conditionally independent given another variable Y , so that the joint distribution of the $n+1$ variables can be decomposed as follows:

$$p(x_1, \dots, x_n, y) = p(y) \prod_i p(x_i|y) \quad . \quad (25)$$

The relations between the variables can be represented with what is known as a *naïve Bayes* network, where Y is the common root node with X_1, \dots, X_n as children. Such a model is amenable for its scarcity of parameters, compared to the full joint distribution, and for the possibility of assessing each $p(X_i|y)$ independently of the others, possibly from different sources of information or at different times.

Having uncertain information about the probability distributions $p(X_i|Y)$ and $p(Y)$ in the form of subjective opinions, we obtain a *naïve Bayes subjective network* (Fig.1).

As an example, let us suppose we want to construct a subjective network for detecting type 2 diabetes (T2D), which will be the common root node, from four of its major risk factors, which will be the children: obesity, old age, family history of T2D, and past episodes of high blood glucose. In this case, the choice of the naïve Bayes network structure is practical: information on T2D prevalence and on the probability distribution of the children nodes with and without diabetes is easy to gather from the appropriate medical sources, whereas information on the joint distribution of the four variables would be much harder to get. Then, in constructing the input opinions $\omega_{X_i|y} = (b_{X_i|y}, u_{X_i|y}, a_{X_i|y})$, the uncertainty mass could be set higher on the conditional opinions when y is *true* and lower when y is *false*, on the account of the much larger amount of samples from which the latter probabilities are probably estimated. Furthermore, uncertainty mass could be set higher on the opinions on family history and past episodes of high blood glucose than on the ones on age and obesity, on the account of the latter two being more reliable to assess precisely and not being based on the memory of past events or on historical clinical records. In the lack of clear evidence for T2D in a particular case, the opinion ω_Y could be set to vacuous, where the only relevant information we use is the domain knowledge (statistics on T2D in the population, for example) reflected in the base rate distribution a_Y .

3.3 Fused Subjective Networks

The initial ideas about subjective networks [Jøsang, 2008] assume a DAG that follows the available input information and a conditional opinion on a child node is provided for each of its parents separately. This is very often more appropriate




Figure 2: A V-structure fused subjective network.

then a Bayesian subjective network for modelling opinions on variables connected in a graph containing V-structures, like the one given in Figure 2. In such cases, it is easier for the analyst to provide the opinions $\omega_{Y|X_1}$ and $\omega_{Y|X_2}$ separately, than the opinions $\omega_{Y|X_1X_2}$ (which is necessary if we want to construct a Bayesian subjective network with the same DAG).

For example, an expert might have an opinion on the probability of military aggression over a country A from a country B conditional on the troop movements in B , and opinion on the probability of military aggression conditional on the political relations between the countries, but is not able to merge these opinions in a single opinion about military aggression conditional on both the factors considered together.

We call a *fused subjective network* a DAG with a set of conditional subjective opinions $\omega_{Y|X}$ associated with an arrow from X to Y , for each arrow in the graph. In addition, the fused subjective networks representation assumes that base rates for the root nodes in the DAG are also available.

Unlike in the case of Bayesian subjective networks, the projected probability distributions of the subjective opinions in a fused subjective network do not necessarily form a Bayesian network with the given DAG, which is the substantial difference between the two representations.

Note that in naïve Bayes networks and, in general, subjective networks with a DAG that is a single-rooted tree, every node has at most one parent, hence, this type of subjective networks is in the intersection of the fused and Bayesian ones.

4 Inference in Subjective Networks

The inference in classical Bayesian networks reduces to the following: given that the values of some of the variables (evidence variables E) are observed, to find the probability of any other variable or set of variables in the network (target variables T), which is to find the conditional probability of the target given an instantiation of the evidence, $p(T|e)$. In subjective networks, the evidence does not necessarily mean an observation. Namely, an analogue to observing the value of a variable, in the case of a subjective opinion is assigning a belief 1 to that value of the variable (based on direct observation, or just a strong opinion), and that gives an evidence in the form of an absolute opinion. In general, we could have evidence in the form of a general type of subjective opinion on E , ω_E , and would like to be able to account for it, i.e. to be able to update the opinion on the target variables conditional on this kind of evidence.

In subjective networks we can distinguish among three different types of subjective evidence:

- *absolute evidence* - evidence in the form of absolute opinion, i.e. instantiation of the evidence variables,

- *dogmatic evidence* - evidence in the form of dogmatic opinion on the evidence variables, and
- *uncertain evidence* - evidence in the form of a subjective opinion with uncertainty greater than zero.

For the derived opinion on the target variable T we will use the notation $\omega_{T||e}$ in the case of absolute evidence, and $\omega_{T||E}$ in the case of dogmatic or uncertain evidence.

Depending on whether E is a set of one or more variables, we can further distinguish between:

- *single evidence* - evidence on one variable only, and
- *multiple evidence* - evidence on more than one variable.

4.1 Inference in a Two-node Network

In this section we briefly summarize the operations of deduction and abduction defined in [Jøsang, 2008; Jøsang and Sambo, 2014] for conditional reasoning with two variables. We assume we have a two-node subjective network, where X is the parent and Y is the child node, and the subjective opinions $\omega_{Y|X} = \{\omega_{Y|x} \mid x \in X\}$ are available, along with the base rate distributions a_X and a_Y .

Subjective Logic Deduction

Given the input opinions $\omega_{Y|X}$ and a subjective evidence ω_X , the goal is to *deduce* a subjective opinion on Y , $\omega_{Y||X}$.⁵

The projected probability distribution of Y in this case will be determined by:

$$P(y||X) = \sum_{x \in X} P(x)P(y|x). \quad (26)$$

For the belief masses of the deduced opinion on Y , we assume the following: The unconditional beliefs of the deduced opinion are at least as large as the minimum of the conditional beliefs:

$$b_{y||X} \geq \min_{x \in X} \{b_{y|x}\}, \quad (27)$$

for every $y \in Y$. This is a natural assumption, which can also be found as a *principle of plausible reasoning* for example in [Pearl, 1990]. Then we first determine the uncertainty mass $u_{Y||\hat{X}}$ corresponding to a vacuous evidence opinion on X , as the maximum possible uncertainty under the conditions in Eq.(27) and Eq.(26). The uncertainty of the deduced opinion from ω_X is then determined as a weighted average:

$$u_{Y||X} = u_X u_{Y||\hat{X}} + \sum_{x \in X} b_x u_{Y|x}. \quad (28)$$

Once we have the uncertainty mass of the deduced opinion, the beliefs are easily derived using the projected probabilities and Eq.(13).

Subjective Logic Abduction

Given the set of opinions $\omega_{Y|X}$ and an opinion ω_Y , the goal is to *abduce* an opinion on X , denoted by $\omega_{X||Y}$.⁶

⁵This can be considered both a Bayesian subjective network with no evidence, or a fused subjective network with a subjective evidence on X .

⁶This case can only be classified as fused subjective network with a subjective evidence on Y .

The operation of abduction first “inverts” the given set of conditional opinions $\omega_{Y|X}$ into a set $\omega_{X|Y}$, and then applies deduction on $\omega_{X|Y}$ and ω_Y to derive the opinion $\omega_{X||Y}$. The projected probability distribution of the inverted opinion $\omega_{X|y}$ is obtained as follows:

$$P(x_j|y_i) = \frac{a_{x_j} P(y_i|x_j)}{\sum_{t=1}^l a_{x_t} P(y_i|x_t)}. \quad (29)$$

Then its uncertainty mass $u_{X|y}$ is obtained by heuristic procedure which takes the maximum possible uncertainty value compatible with Eq.(29) and adjusts it using the average uncertainty of the input conditionals $\omega_{Y|X}$ and the irrelevance of X to the value y , for details see [Jøsang and Sambo, 2014]. The beliefs $b_{x|y}$ are again consequence of the projected probabilities and the uncertainty.

4.2 Inference in Bayesian Subjective Networks

Let us assume that we are given a Bayesian subjective network S_n and absolute evidence ω_E^e on the set of variables E . Given the instantiation e of E , we want to find the conditional subjective opinion on a target T , $\omega_{T||e}$. We assume the following:

1. The projected probability of the derived opinion is determined from the projected Bayesian network in a classical way, i.e. $P(T|e)$ is determined in $P(S_n)$ using the standard Bayesian networks reasoning methods.
2. All the conditional and marginal base rate distributions in the subjective network are either given *a priori*, or determined from the ones provided in the initial opinions in the network by Bayesian reasoning.
3. The uncertainty and the beliefs of the derived opinion satisfy certain constraints, like, for example, some of the conditions given in Eq.(24) and Eq.(27).

The first assumption provides a way of determining the projected probability of the derived opinion and is a starting point in deriving the opinion $\omega_{T||e}$. Namely, having determined the projected probability of an opinion and considering the second assumption above, we have obtained by Eq.(13) a system of m linear equations with the beliefs and the uncertainty mass of the derived opinion as variables, where m is the cardinality of the target T . We obtain one more equation over the same variables from the additivity property for the beliefs and uncertainty of subjective opinions, given in Eq.(12). This means that we have a system of $m+1$ equation with $m+1$ variables, which might seem to fully determine the required opinion $\omega_{T||e}$. However, the projected probabilities on the left-hand side of the equations in Eq.(13) also add up to 1, which makes this system dependent. Hence, the system has an infinite number of solutions, i.e. there are infinitely many subjective opinions with the derived projected probability, and we have to apply additional constraints on beliefs and the uncertainty mass (assumption 3.) to choose a single one as a solution.

The above discussion implies the following: If we find a suitable way of determining the uncertainty mass of the derived opinion, the beliefs follow from Eq.(13) (the base rate is either a priori given or determined from the given base rates),

and the opinion is fully derived. While this is successfully applied in the deduction for two variables described in the previous section, in general, it remains a challenge to provide a meaningful way of propagating the given uncertainty masses throughout the network in a way that would give reasonable belief mass values (that satisfy the initially set constraints) as a consequence. Also, there might not exist a unique way of propagating the uncertainty, and how we decide to do it can be context-dependent.

The above described inference procedure would operate over multinomial opinions. It would be possible though to provide input information in a form of hyper opinions, in which case their multinomial approximations (described at the end of Section 2.2) can be used in the inference procedure, to derive a multinomial opinion on the target. This is an advantage in some sense, since one usually has the input information in the more vague, hyper opinion form, and wants to have the output as a multinomial opinion, i.e. have the distribution of beliefs over the values rather than set of values.

The inference from dogmatic or uncertain evidence remains a challenge, for in that case we can not have the assumption 1, namely: instantiating the evidence variables E in a given subjective network S_n with a subjective opinion ω_E that is not absolute, we simultaneously provide a new projected probability distribution of E , which, in general, differs from the one that would be derived by Bayesian reasoning in $P(S_n)$.

4.3 Inference in Fused Subjective Networks

Let us assume we have a fused subjective network with a singly-connected DAG (only one path between any two nodes) and evidence given in the form of an (unconditional) opinion on an arbitrary node in the network X . The goal of the inference is to determine (unconditional) opinion on Y based on the given evidence, namely $\omega_{Y||X}$. This can be done by propagating the evidence through the path from X to Y and performing a series of deduction/abduction operations depending on the direction of the arrow at each step.

The inference from multiple evidence in the case we have a subjective opinion on each of the sources (evidence variables) separately, can be done by deriving one subjective opinion on the target for each of the evidence variables, and then fusing the resulting opinions. The operation of fusion [Jøsang and Hankin, 2012] can be applied for this purpose and there is a variety of existing fusion operators that take in account the degree of independence of the sources of evidence.

In some cases, inference in a fused subjective network can be done by first transforming it into a Bayesian one in the following way: for every V-structure with parents X_1, \dots, X_n and child Y , we invert the given set of conditionals opinions $\omega_{Y|X_i}$, $i = 1, \dots, n$ into $\omega_{X_i|Y}$ as described in the abduction operation in Section 4.1. This means that we invert the V-structure into a naïve Bayes network where Y is a parent of X_1, \dots, X_n . Because of the conditional independences in the naïve Bayes, we can apply the product operation from [Jøsang and McAnally, 2004] on the opinions $\omega_{X_i|y}$, $i = 1, \dots, n$ to obtain the opinion $\omega_{X_1 \dots X_n|y}$, for every $y \in \mathbb{Y}$, i.e. the set of opinions $\omega_{X_1 \dots X_n|Y}$. At the end, we invert again to obtain the set $\omega_{Y|X_1 \dots X_n}$.

5 Inference through the Dirichlet Representation

This section provides an alternative approach towards inference in subjective networks, based on the Dirichlet pdf representation of subjective opinions introduced in Section ??.

In a subjective network, evidence has been collected to form subjective opinions about the conditional probabilities. In other words, each conditional probability distribution $p(X_i|pa(X_i))$ is represented as a k_i -dimensional Dirichlet distributed random variable, where $k_i = |\mathbb{X}_i|$. Because of the Markov property, these Dirichlet distributed random variables are also statistically independent.

One important goal of inference from absolute evidence in subjective networks is to derive an opinion $\omega_{X_i||e}$ for a given instantiation e of evidence variables E , and a single target node X_i not in E . In terms of the Dirichlet representation, determining $\omega_{X_i||e}$ is equivalent to determining the appropriate Dirichlet pdf to represent the uncertainty about the probability distribution $p(X_i|e)$. According to Section 2.1, this probability distribution is expressed through the input probability distributions in the graph in the following way:

$$p(x_i|e) = \frac{\sum_{X_j \in V \setminus E \cup \{X_i\}} \prod_{k=1}^n p(x_k|pa(X_k))}{\sum_{X_j \in V \setminus E} \prod_{k=1}^n p(x_k|pa(X_k))}, \quad (30)$$

where $pa(X_k)$ is the instantiation of the parents of X_k that corresponds to x_i and e .

For a standard Bayesian network, the execution of Eq.(30) can be accomplished as a series of variable elimination steps [Zhang and Poole, 1994]. For subjective networks, the probability distributions involved in the right-hand side of Eq.(30) are Dirichlet distributed random variables and exact inference becomes more challenging. The target probability distribution $p(X_i|e)$ is a k_i -dimensional random variable characterized through the independent Dirichlet distributed random variables $p(X_k|pa(X_k))$. Through a change of variables process, it is possible to determine the actual pdf for $p(X_i|e)$, which in general is not a Dirichlet pdf.

In order to infer a subjective opinion on X_i given e by means of the transformation in Eq.(19), we need to approximate this pdf by a Dirichlet pdf. We choose to use a moment matching approach to determine the best Dirichlet pdf to approximate the pdf of $p(X_i|e)$. First, the mean value of this Dirichlet pdf, $m(X_i|e)$, must equal the expectation of the actual pdf for $p(X_i|e)$. Then, the Dirichlet strength s is selected so that the second order moments of the actual target distribution matches that of the Dirichlet distribution as much as possible in the mean squared sense. The matching of the second order moments is perfect only for binary variables. The moment matching method to determine the Dirichlet strength has been implemented for partial observation updates and deduction in [Kaplan *et al.*, 2015a] and [Kaplan *et al.*, 2013], respectively. In the general case where the evidence can come from the descendants of X_i , a closed form solution for the first and second order expectation of Eq.(30) does not exist because of its fractional form, and one must resort to numerical integration over N Dirichlet distributed random variables, where N is the number of input probability distributions in

the network. Such a moment matching method is only computationally feasible for the smallest of networks.

Current research is looking at extending the sum-product algorithm [Wainwright and Jordan, 2008]. Such an approach develops a divide and conquer strategy that will provide means to propagate one piece of evidence at a time. Then the effects of an observation coming from the antecedents is propagated forward via subjective logic deduction (as in [Kaplan *et al.*, 2013]), and a backwards process will enable the computation of the effect of an observation coming from a descendant node. At each stage in the process, the stored conditionals are approximated by Dirichlet distributions using the moment matching method. Finally, the inference of the target opinion from combined evidence is accomplished by normalizing the opinions conditional on evidence coming from different directions. The first steps of this normalization process has been studied in [Kaplan *et al.*, 2015b] for the case of a three-node chain of binary variables.

The evaluation of forward/backward propagation along with normalization over chains is the next step. The intermediate results will be stored as subjective opinions, which means that the inference via normalization will only be an approximation of moment matching of Eq.(30), which is not making any Dirichlet approximation about the marginal distribution for the intermediate nodes between the evidence E and X_j . This is in contrast to the sum-product algorithm over Bayesian networks, which provides exact inference. The plan is to study computational efficiency and accuracy of imposing the Dirichlet approximation as the effects of the observations propagate over the “uncertain” probabilistic edges.

The development and evaluation of inference techniques over subjective networks will consider increasing complexity in various dimensions. One dimension is the topology of the network, where we will first study chains and then expand to trees and eventually arbitrary DAGs where we will need to modify the sum-product framework. Another dimension is the complexity of the subjective opinions: We start with binary ($k_i = 2$) and multinomial opinions ($k_i > 2$), to finally consider hyper opinions ($2^{k_i} - 2$). The quality of the observations over E provides another complexity dimension to explore. Initially, we will only consider inference from absolute opinions, which are equivalent to instantiation of variables, but in future work we plan to consider inference from general type of opinions.

6 Conclusions and Future Work

We introduced subjective networks as structures for conditional reasoning with uncertain probabilistic information, represented in the form of subjective opinions on random variables. In this way both the input information and the inferred conclusions in the modelled scenario incorporate a current analysis of beliefs and domain knowledge, at the same time taking the uncertainty about the probabilities explicitly into account.

The discussed inference problems in subjective networks lead to the following inference approaches to be studied in future work: global uncertainty propagation in a Bayesian subjective network, piece-wise inference in fused networks,

and a statistical moment matching approach.

References

- [Gelman and others, 2004] A. Gelman et al. *Bayesian Data Analysis, 2nd ed.* Chapman and Hall/CRC, Florida, USA, 2004.
- [Jøsang and Hankin, 2012] Audun Jøsang and Robin Hankin. Interpretation and Fusion of Hyper Opinions in Subjective Logic. In *Proceedings of the 15th International Conference on Information Fusion (FUSION 2012)*, Singapore, July 2012.
- [Jøsang and McAnally, 2004] A. Jøsang and D. McAnally. Multiplication and Comultiplication of Beliefs. *International Journal of Approximate Reasoning*, 38(1):19–51, 2004.
- [Jøsang and Sambo, 2014] Audun Jøsang and Francesco Sambo. Inverting Conditional Opinions in Subjective Logic. In *Proceedings of the 20th International Conference on Soft Computing (MENDEL 2014)*, Brno, 2014.
- [Jøsang, 2001] A. Jøsang. A Logic for Uncertain Probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 9(3):279–311, June 2001.
- [Jøsang, 2008] A. Jøsang. Conditional Reasoning with Subjective Logic. *Journal of Multiple-Valued Logic and Soft Computing*, 15(1):5–38, 2008.
- [Kaplan *et al.*, 2013] Lance M. Kaplan, Murat Sensoy, Yuqing Tang, Supriyo Chakraborty, Chatschik Bisdikian, and Geeth de Mel. Reasoning under uncertainty: Variations of subjective logic deduction. In *Proc. of the International Conference on Information Fusion*, Istanbul, Turkey, July 2013.
- [Kaplan *et al.*, 2015a] Lance Kaplan, Murat Şensoy, Supriyo Chakraborty, and Geeth de Mel. Partial observable update for subjective logic and its application for trust estimation. *Information Fusion*, 2015.
- [Kaplan *et al.*, 2015b] Lance Kaplan, Magdalena Ivanovska, Audun Jøsang, and Francesco Sambo. Towards subjective networks: Extending conditional reasoning in subjective logic. To appear in the Proc. of the International Conference on Information Fusion, 2015.
- [Pearl, 1988] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman Publishers, 1988.
- [Pearl, 1990] Judea Pearl. Reasoning with belief functions: An analysis of compatibility. *International Journal of Approximate Reasoning*, 4(5):363 – 389, 1990.
- [Shafer, 1976] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [Wainwright and Jordan, 2008] Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1):1–305, January 2008.
- [Zhang and Poole, 1994] Nevin L Zhang and David Poole. A simple approach to bayesian network computations. In *Proc. of the Tenth Canadian Conference on Artificial Intelligence*, 1994.