Uncertainty-aware Situational Understanding

Richard Tomsett[a], Lance Kaplan[b], Federico Cerutti[c], Paul Sullivan[d], Daniel Vente[e], Marc Roig Vilamala[c], Angelika Kimmig[c], Alun Preece[c], and Murat Şensoy[e]

[a]IBM Research, UK;
[b]ARL, USA;
[c]Cardiff University, UK;
[d]Intelpoint Inc., USA;
[e]Ozyegin University, Turkey

ABSTRACT

Situational understanding is impossible without causal reasoning and reasoning under and about uncertainty, i.e. probabilistic reasoning and reasoning about the confidence in the uncertainty assessment. We therefore consider the case of subjective (uncertain) Bayesian networks. In previous work we notice that when observations are out of the ordinary, confidence decreases because the relevant training data—effective instantiations—to determine the probabilities for unobserved variables—on the basis of the observed variables—is significantly smaller than the size of the training data—the total number of instantiations. It is therefore of primary importance for the ultimate goal of situational understanding to be able to efficiently determine the reasoning paths that lead to low confidence whenever and wherever it occurs: this can guide specific data collection exercises to reduce such an uncertainty. We propose three methods to this end, and we evaluate them on the basis of a case-study developed in collaboration with professional intelligence analysts.

Keywords: Uncertainty aware, situational understanding

1. INTRODUCTION

Situational understanding is impossible without causal reasoning and reasoning under and about uncertainty. Indeed its traditional definition[1] situational understanding is the "product of applying analysis and judgement to the unit's situation awareness to determine the relationships of the factors present, and form logical conclusions concerning threats to the mission accomplishment, opportunities for mission accomplishment, and gaps in information." The UK Ministry of Defence Doctrine[2] goes beyond and explicitly mention that (situational) "Understanding involves acquiring and developing knowledge to a level that enables us to know why something has happened or is happening (insight) and be able to identify and anticipate what may happen (foresight)."

As we already argued,[3] Artificial Intelligence (AI) holds the promise to provide efficient and effective methods for supporting humans in situational understanding in a human/machine collaborative effort. The recent successes of machine learning (ML)[4,5] have sparked new interest in artificial intelligence (AI)&ML systems.[6] For the most part, the emphasis is on improving the accuracy of the probabilistic inferences machine learning: this is *aleatoric uncertainty*.

We can indeed distinguish between two types of uncertainty. Let us consider a coin, and let us toss it three times. Let us imagine we obtain head two times: from a frequentist point of view, the probability of having head is $\frac{2}{3}$. This is *aleatoric uncertainty*: assuming that we have enough statistically significant samples of a phenomenon, we can then assess the probability of it. However, this raises the question: "how confident are we that the coin is not fair"? We name this level of confidence *epistemic uncertainty*.

Because of the stringent need of situational understanding to express causal relationships we focus on an uncertainty-aware—i.e. able to consider both aleatoric and epistemic uncertainty—Bayesian network approach, Subjective Bayesian Network[7] that we summarise in Section 3. In previous work[8] we notice that when observations are out of the ordinary, epistemic uncertainty increases because the relevant training data—effective instantiations—to determine the probabilities for unobserved variables—on the basis of the observed variables—is significantly smaller than the size of the training data—the total number of instantiations. It is therefore of primary importance for the ultimate goal of situational understanding to be able to efficiently determine the reasoning paths that lead to high epistemic uncertainty whenever and wherever it occurs: this can guide specific data collection exercises to reduce such an uncertainty.

The main contribution of this paper is a case-study analysis of techniques for assess causal links strength in SBN (Section 4). In particular we derive an approach based on mutual information between two random variables originally proposed in the context of Bayesian networks,[9] and we compare it with two native methods exploiting the uniqueness of the SBN approach. We then discuss an articulated empirical analysis (Section 5) that shows strengths and weaknesses of our three methods for analysing the strengths of the causal links in the network.

2. MOTIVATIONAL SCENARIO

To demonstrate the need for uncertainty-aware AI&ML, let us recall a scenario[8] where we consider a route planning operation within a purposely simplified peace keeping scenario to highlight the salient value of uncertainty-awareness. In this scenario, the GoodGuys are trying to stabilize a city. Consider that a commander of the GoodGuys needs to transport some medical supplies from headquarters to a local hospital. There are three routes through a city populated by two factions, the Capulets and the Montagues. These factions have a tumultuous relationships with each other and the GoodGuys. One route (Route A) goes through the center of the Capulet territory, and another route (Route C) goes through the center of the Montague territory. The remaining route (Route B) skirts the border of Capulet and Montague territory and passes by the central marketplace.

There is a surveillance camera at the marketplace that is connected to a neural network that classifies between normal and violent activity. For the purpose of this work, we assume that the neural network is able to output a Dirichlet distribution over the set of all the labels it has been trained on, e.g. implementing the Evidential Deep Learning.[10] The commander also has information about the number of people who visit the community centers every day within the Capulet and Montague territories. Given these observations, the commander wants to decide what route to use and the level of armored escorts needed to protect the transport operation. The commander is using the Bayesian network model illustrated in Figure 1 to predict the route conditions. This model indicates that the daily dispositions of each faction towards the GoodGuys influences the attendance at the civic centers, the marketplace activity, and the safety of the three routes.

3. UNCERTAINTY-AWARE SOLUTIONS FOR EFFECTIVE SITUATIONAL UNDERSTANDING

In the era of big data, many assume that the parameters of an AI&ML system can be interpreted as point estimates and the prediction accuracy of the system can generalize as conditions change. However, for the scenarios described above, training data is sparse and the situations are highly dynamic. As a result, the parameters are highly uncertain and possibly changing over time. This uncertainty propagates through the process to infer the beliefs in the state values of the latent variables crucial for decision making. When the conditions that the AI&ML system operates under become vastly different than the training conditions, these uncertainties can be high indicating the inferencing power of the system is currently poor. This section reviews current work to understand how to extract these uncertainties.

3.1 Introduction to Subjective Logic

When probabilities are uncertain—for instance because of limited observations—such an uncertainty can be captured by a Dirichlet distribution, or by a Beta distribution, if there are only two mutually exclusive outcomes.

Let us consider a random variable $\boldsymbol{X}$ that can take on a finite set of values over domain $\mathbb{X} = \langle X_1, \ldots, X_K \rangle$. The value of $\boldsymbol{X}$ does change over different instantiations, and there is an underlying ground truth value for the appearance probabilities $\mathbf{p}_X$. Therefore, the value for $\boldsymbol{X}$ at a given instantiation is drawn from a multinomial distribution. After tabulating the counts of the value of $\boldsymbol{X}$ over N_{ins} instantiations, one can represent knowledge of $\mathbf{p}$ as following a Dirichlet distribution. This is because the Dirichlet is the conjugate-prior for the multinomial distribution.

The Dirichlet probability density distribution is

$$D(\mathbf{p}|\boldsymbol{\alpha}) = \begin{cases} \frac{1}{B(\alpha)} \prod_{i=1}^{K} p_i^{\alpha_i - 1} & \text{for } \mathbf{p} \in \mathcal{S}_K, \\ 0 & \text{otherwise,} \end{cases}$$

where the parameters $\boldsymbol{\alpha}_X = \langle \alpha_1, \cdots, \alpha_K \rangle$, $\mathcal{S}_K$ is the K-dimensional unit simplex, i.e.,

$$\mathcal{S}_K = \left\{ \mathbf{p} = \langle p_1, \ldots, p_K \rangle \;\middle|\; \sum_{i=1}^{K} p_i = 1 \text{ and } 0 \leqslant p_1, \ldots, p_K \leqslant 1 \right\},$$

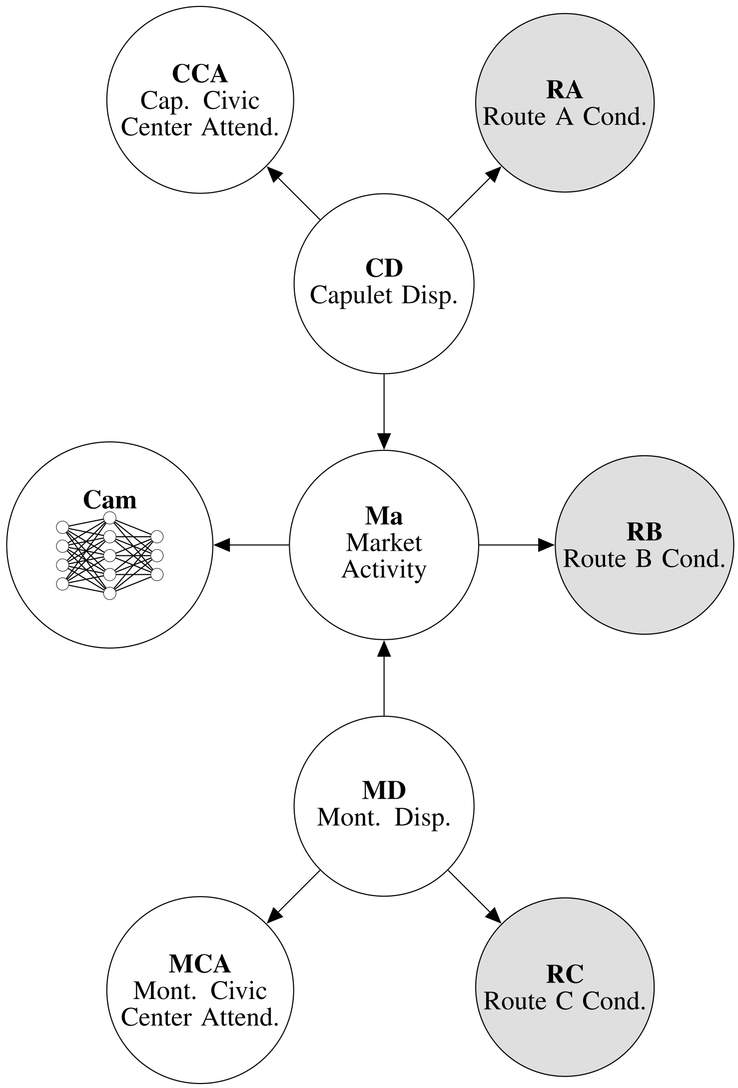

Figure 1: Bayesian network model describing the probabilistic interactions of the various variables that influences how to infer route safety in light of the observations including the civic center engagement and marketplace activity likelihood via the surveillance camera.

and $B(\boldsymbol{\alpha})$ is the K-dimensional multinomial beta function.[11] Note that when $\alpha_i = 1$ for $i = 1, \ldots, K$, the distribution is uniform over the K-dimensional simplex. This represents a non-informative prior. With such a prior

$$\mathbf{e_X} = \boldsymbol{\alpha}_X - \mathbf{1} \tag{1}$$

represents the counts of variable values that leads to posterior knowledge of the $\mathbf{p}_X$ by the Dirichlet distribution with parameters $\boldsymbol{\alpha}_X$. Thus, $\mathbf{e_X}$ represents the evidence about the generating probability $\mathbf{p}_X$ of the values of $\boldsymbol{X}$.

Given a Dirichlet-distributed knowledge about the random variable X,

$$s_{\boldsymbol{X}} = \sum_{i=1}^{K} \alpha_i \tag{2}$$

is its *Dirichlet strength* and

$$\mu_{\boldsymbol{X}}(X_i) = \frac{\alpha_i}{s_{\boldsymbol{X}}} \tag{3}$$

is the mean appearance probability that $\boldsymbol{X} = X_i$. From (2) and (3) the beta parameters can equivalently be written as:

$$\boldsymbol{\alpha}_X = \langle \mu_{\boldsymbol{X}}(X_1) s_{\boldsymbol{X}}, \ldots, \mu_{\boldsymbol{X}}(X_K) s_{\boldsymbol{X}} \rangle. \tag{4}$$

The variance for the appearance probabilities of the random vector $\boldsymbol{X}$ is

$$\sigma^2_{\boldsymbol{X}}(X_i) = \frac{\mu_{\boldsymbol{X}}(X_i)(1 - \mu_{\boldsymbol{X}}(X_i))}{s_{\boldsymbol{X}} + 1}. \tag{5}$$

3.1.1 Subjective Logic

Subjective logic[12] provides an alternative, more intuitive, way of representing the parameters of a Dirichlet-distributed random vector. A subjective opinion about $\boldsymbol{X}$ is a tuple $\omega_X = \langle b_{X_1}, \ldots, b_{X_K}, u_X, a_X \rangle$, representing the beliefs, uncertainty, and prior probability about $\boldsymbol{X}$. These values are non-negative and $\sum_{i=1}^{K} b_{X_i} + u_X = 1$. The projected probability $P(x_i) = b_{X_i} + u_X \cdot a_{X_i}$, provides an estimate of the ground truth probability p_{x_i}.

The mapping from a Dirichlet-distributed appearance probability for the random value $\boldsymbol{X}$ with parameters $\boldsymbol{\alpha}_X$ to a subjective opinion is:

$$\omega_X = \left\langle \frac{\alpha_1 - W a_{X_1}}{s_X}, \ldots, \frac{\alpha_K - W a_{X_K}}{s_X}, \frac{W}{s_X}, a_X \right\rangle \tag{6}$$

With this transformation, the mean of X is equivalent to the projected probability $P(x)$, and the Dirichlet strength is inversely proportional to the uncertainty of the opinion:

$$\mu_X(x_i) = P_{\boldsymbol{X}}(x_i) = b_{X_i} + u_X a_{X_i}, \quad s_X = \frac{W}{u_X}, \tag{7}$$

for all $x \in \mathbb{X}$.

Conversely, a subjective opinion ω_X translates directly into a Beta-distributed random variable with:

$$\boldsymbol{\alpha}_X = \left\langle \frac{W}{u_X} b_{X_1} + W a_{X_1}, \ldots, \frac{W}{u_X} b_{X_K} + W a_{X_K} \right\rangle. \tag{8}$$

Given that in the absence of any evidence, the prior probability for $\mathbf{p_X}$ is $\boldsymbol{\alpha}_X = W a_X$, the evidence for opinion ω_X is

$$\boldsymbol{e}_X = \left\langle \frac{W}{u_X} b_{X_1}, \ldots, \frac{W}{u_X} b_{X_K} \right\rangle. \tag{9}$$

In this paper, $W = 2$ and $a_{X_i} = 1/K$ for $i = 1, \ldots, K$.

Subjective logic provides a set of rules to infer the effective Dirichlet strength and projected probabilities of latent variables from the logical entailment with variables for which opinions come from direct observations. For instance, subjective opinions naturally extend to subjective conditional opinions, where for example, the opinion for X conditioned on Y and Z is interpreted as the set $\{\omega_{X|y,z} : y \in \mathbb{Y}, z \in \mathbb{Z}\}$, and $\omega_{X|y,z}$ represents the effective number of times that $X = x$ for $x \in \mathbb{X}$ when $Y = y$ and $Z = z$ while jointly observing X, Y, and Z. Overall, the efficacy of subjective logic reasoning can be evaluated via simulations to determine 1) how accurate the projected probabilities of the inferred opinions match the ground truth probabilities, and 2) how accurately the uncertainty (or Dirichlet strength) of the inferred opinions represents the 'spread' between the projected and ground truth probabilities.

3.1.2 Consistency of Opinions

This section develops a likelihood ratio to determine whether or not two subjective opinions represents the same variable, i.e., the appearance probabilities are equivalent. Let e_k^0 and e_k^1 for $k = 1, .., K$ represent the evidence for two opinions ω_{X_0} and ω_{X_1}, respectively.The question is whether or not $\mathbf{p}_{X_0} = \mathbf{p}_{X_1}$ or equivalently $\boldsymbol{X}_0 = \boldsymbol{X}_1$.

Let's say that the probabilities for the results of obtaining $\mathbf{e}^0$ is p_k for $k = 1, \ldots, K$. Then the likelihood for forming the evidence is simply

$$l = \prod_{k=1}^{K} p_k^{e_k^0}. \tag{10}$$

Let H_1 be the hypothesis that the appearance probabilities are equivalent, i.e., $\boldsymbol{X}_0 = \boldsymbol{X}_1$. Then if $\mathbf{p}$ is known, the l would be the likelihood that H_1 is true. However, $\mathbf{p}$ is only known within a Dirichlet distribution based on the evidence $\mathbf{e}^1$, and one must instead consider the expected likelihood over possible appearance probabilities for $\boldsymbol{X}_1$ as

$$\begin{aligned} \ell_1 &= \int_{S_K} \prod_{k=1}^{K} p_k^{e_k^0} D(\mathbf{p}|\mathbf{e}^1 + \mathbf{1}) d\mathbf{p}, & (11) \\ &= \int_{S_K} \prod_{k=1}^{K} p_k^{e_k^0} \frac{1}{B(\mathbf{e}^1 + \mathbf{1})} \prod_{k=1}^{K} p_k^{e_k^1} d\mathbf{p}, & (12) \\ &= \frac{B(\mathbf{e}^0 + \mathbf{e}^1 + \mathbf{1})}{B(\mathbf{e}^1 + \mathbf{1})}. & (13) \end{aligned}$$

The H_0 hypothesis is $\boldsymbol{X}_0 \neq \boldsymbol{X}_1$. Then we do not know anything about the appearance probability that generated the evidence $\mathbf{e}^0$. That is, the distribution for the p's are uniform (a non-informative prior) and the expected likelihood is

$$\begin{aligned} \ell_0 &= \frac{1}{B(\mathbf{1})} \int_{S_K} \prod_{k=1}^{K} p_k^{e_k^0} dp & (14) \\ &= (K-1)! B(\mathbf{e}^0 + \mathbf{1}). & (15) \end{aligned}$$

Note that $B(\mathbf{1}) = \frac{1}{(K-1)!}$.

For the two opinions, the ratio of the expected likelihood that the two opinions are the same over that they are different is

$$\Lambda(\omega_{X_0}, \omega_{X_1}) = \frac{B(\mathbf{e}^0 + \mathbf{e}^1 + \mathbf{1})}{B(\mathbf{e}^0 + \mathbf{1}) B(\mathbf{e}^1 + \mathbf{1})} \frac{1}{(K-1)!}. \tag{16}$$

This ratio of expected likelihoods is large when it is likely the two opinions represent the same variable. On the hand, the expected likelihood ratio is small (near zero) when it is likely the two opinions represent different variables. The likelihood ratio is near one when it is difficult to discern if the variables are the same or different (due to a lack of evidence).

3.2 Subjective Logic Bayesian Network

Subjective Bayesian networks are uncertain Bayesian networks where the uncertain knowledge of the conditional probabilities is represented as Dirichlet distributions, i.e., subjective opinions. As a result, the probability of any unobserved variable value conditioned on the evidence, i.e., the values of the observed variables, is only known within a distribution. It is possible to approximate that distribution as a subjective opinion. In,[13] *subjective belief propagation* extends belief propagation[14] so that the π- and λ-messages, which are passed from parents and children variables, are treated as subjective opinions rather than as point values. As such, they are characterized by a mean probability and Dirichlet strength.

The SBP formulation approximates output messages as Dirichlet distributed using the methods of moments and a first order Taylor series approximation to determine the mean and variance of the output messages in light of Dirichlet distributed input messages. The details of the derivations are provided in[13] for a singularly connected network of binary variables so that the variables and messages simplify to beta distributions. This subsection summarizes the operations to implement SBP.

At its core, SBP requires three operation: 1) Forward propagation, 2) backward propagation and 3) fusion. The forward propagation combines opinions of the π-messages from a node's parents with the node's conditional opinion to form a single π-opinion about the node. Given a node X with m parents U_i for $i = 1, \ldots, m$, the subjective opinions of the π-messages sent to X are characterized by the mean probabilities $\mu_{\pi_{U_i,X}}(x)$ and Dirichlet strengths $s_{\pi_{U_i,X}}$. Node X forms its π-opinion by forward propagating the opinions for π-messages from its parents in light of its conditional opinion, which we express as

$$\boldsymbol{\omega}_{\pi_X} = \boldsymbol{\omega}_{P_{X|U_1,\ldots,U_m}} \circledcirc \boldsymbol{\omega}_{\pi_{U_1,X}}, \ldots, \boldsymbol{\omega}_{\pi_{U_m,X}}. \tag{17}$$

This operation determines the means and Dirichlet strength for the node's π-opinion from the respective means and strengths of the π-messages and the conditionals via

$$\mu_{\pi_X}(x) = \sum_{u_1,\ldots,u_m} \mu_P(x|u_1,\ldots,u_m) \prod_{i=1}^{m} \mu_{\pi_{U_i,X}}(u_i) \quad \text{and} \quad s_{\pi_X} = \frac{\mu_{\pi_X}(x)(1-\mu_{\pi_X}(x))}{\sigma^2_{\pi_X}} - 1 \tag{18}$$

for $x \in \mathbb{X}$, where the variance is $\sigma^2_{\pi_X} = V_{\pi_X}(x) - \mu^2_{\pi_X}(x)$,

$$V_{\pi_X}(x) = \sum_{u_1,\ldots,u_m} \sum_{u'_1,\ldots,u'_m} g(x,x;u_1,\ldots,u_m;u'_1,\ldots,u'_m) \prod_{i=1}^{m} h(u_i,u'_i), \tag{19}$$

$$g(x,x';u_1,\ldots,u_m;u'_1,\ldots,u'_m) = \mu_{p_{x|u_1\ldots u_m}} \mu_{p_{x|u'_1\ldots u'_m}} + (-1)^{x\neq x'} \delta_{\mathbf{u},\mathbf{u}'} \frac{\mu_{p_{x|u_1\ldots u_m}}(1-\mu_{p_{x|u_1\ldots u_m}})}{s_{X|u_1\ldots u_m}+1}, \quad \text{and} \tag{20}$$

$$h_\pi(u_i,u'_i) = \mu_{\pi_{U_i,X}}(u_i) \mu\pi_{U_i,X}(u'_i) + (-1)^{u_i\neq u'_i} \frac{\mu_{\pi_{U_i,X}}(u_i)(1-\mu_{\pi_{U_i,X}}(u_i))}{s_{\pi_{U_i,X}}+1}. \tag{21}$$

The fusion operation is used to determine the opinion for the node variables in light of the observed evidence and the opinions for the π-messages to send to the children. Given n-opinions ω_{X_i} for $i = 1, \ldots, n$, the fused opinion is $\omega_{X_f} = \bigodot_{i=1}^{n} \omega_{X_i}$. The fusion operation determines the mean and Dirichlet strength for X_f from those of the $X'_i s$ via

$$\mu_{X_f}(x) = \alpha_f \prod_{i=1}^{n} \mu_{X_i}(x) \quad \text{and} \quad s_{X_f} = \left(\sum_{i=1}^{n} \frac{\mu_{X_f}(x)(1-\mu_{X_f}(x))}{\mu_{X_i}(x)(1-\mu_{X_i}(x))} \frac{1}{s_{X_i}+1} \right)^{-1} - 1, \tag{22}$$

where

$$\alpha_f = \sum_{x\in\mathbb{X}} \prod_{i=1}^{n} \mu_{X_i}(x)$$

is a normalizing constant so that the means sum to one.

Now, given that X has k children Y_j for $j = 1, \ldots, k$, the subjective opinions of the λ-messages sent to X are characterized by the mean probabilities $\mu_{\lambda_{U_i,X}}(x)$ and Dirichlet strengths $s_{\lambda_{U_i,X}}$. Node X forms its λ opinion by fusing the opinions of the children λ-messages via

$$\boldsymbol{\omega}_{\lambda_X} = \bigodot_{j=1}^{k} \boldsymbol{\omega}_{\lambda_{Y_j,X}}. \tag{23}$$

The π and λ-opinions are fused to determine the marginal opinion $\boldsymbol{\omega}_{X|o}$ for node X, i.e.,

$$\boldsymbol{\omega}_{X|o} = \boldsymbol{\omega}_{\pi_X} \odot \boldsymbol{\omega}_{\lambda_X}. \tag{24}$$

The π-messages that node X sends to it children Y_j are

$$\boldsymbol{\omega}_{\pi_{X,Y_j}} = \left(\bigodot_{i\neq j} \boldsymbol{\omega}_{\lambda_{Y_i,X}} \right) \odot \boldsymbol{\omega}_{\pi_X} \tag{25}$$

for $j = 1, \ldots, k$.

The opinion $\omega_{\lambda_{X,U_i}}$ for the message that node X sends to parent U_i is determined by the backward propagation operator expressed as

$$\omega_{\lambda_{X,U_i}}(u_i) = \omega_{\lambda_X} \circledcirc_i \omega_{P_{X|U_1,\ldots,U_m}}, \omega_{\pi_{U_1,X}}, \cdots, \omega_{\pi_{U_m,X}}. \tag{26}$$

The operation determines the mean and Dirichlet strength of the outgoing λ-message from those of the conditional and λ-opinion of X and the other π-messages sent to X via

$$\mu_{\lambda_{X,U_i}}(u_i) = \alpha_b \sum_x \lambda_X(x) \sum_{pa(x)\setminus u_i} \mu_P(x|u_1,\ldots,u_i,\ldots,u_m) \prod_{j\neq i} \mu_{\pi_{U_j,X}}(u_j) \quad \text{and} \quad s_{\lambda_{X,U_i}} = \frac{\mu_{\lambda_{X,U_i}}(u_i)(1-\mu_{\lambda_{X,U_i}}(u_i))}{\sigma^2_{\lambda_{X,U_i}}} - 1, \tag{27}$$

where

$$\sigma^2_{\lambda_{X,U_i}} = \alpha_b^2 \left(\mu^2 \lambda_{X,U_i}(\bar{x}) \sigma^2_{uu} + \mu^2 \lambda_{X,U_i}(x) \sigma^2_{\bar{u}\bar{u}} - 2\mu_{\lambda_{X,U_i}}(x) \mu_{\lambda_{X,U_i}}(\bar{x}) \sigma^2_{u\bar{u}} \right), \tag{28}$$

$$\sigma^2_{zv} = \sum_x \sum_{x'} h_\lambda(x,x') \sum_{pa(x)\setminus z} \sum_{pa(x)'\setminus v} g(x,x';u_1,\ldots,z,\ldots,u_m;u'_1,\ldots,v,\ldots,u'_m) \prod_{j\neq i} h_\pi(u_j,u'_j), \tag{29}$$

$$h_\lambda(x,x') = \mu_{\lambda_X}(x)\mu_{\lambda_X}(x') + (-1)^{x\neq x'} \frac{\mu_{\lambda_X}(x)(1-\mu_{\lambda_X}(x))}{s_{\lambda_X}+1}, \tag{30}$$

and α_b is a normalizing constant.

The equations for the mean probability updates in SBP mirror the update equations in standard belief propagation due to the first-order Taylor approximation. In short, SBP provides the same answer as belief propagation in the mean value. The difference is that SBP also provides a quantification of the uncertainty through the Dirichlet strength. On a technical note, SBP will actually increase the Dirichlet strength as computed in the update equations to ensure that all belief values are non-negative. We refer the interested reader to[13] for more details. Finally, the information flow in SBP is exactly the same as in belief propagation. Namely, a node can send a message to one particular neighbor once it receives messages from all of its other neighbors.

In,[13] it is shown that the projected probabilities of the inferred opinions accurately reflect the probabilities that could be inferred if the ground truth conditional probabilities were known. More importantly for uncertainty-awareness, the uncertainty of the inferred opinions faithfully characterized the accuracy of the projected probabilities. Specifically, empirical results demonstrated that the divergence between the desired and actual confidence bound significance is low. Current and future work is expanding the inference to efficiently work for general (non-binary) variables under general (possibly loopy) network structures.

The network above can only consider binary observations. The observation from a camera in Figure 1 is a raw video sequence that clearly is not binary. One way to accommodate such a more complex variable is to assume that backpropagation is accomplished by a neural network. In other words, the probabilities reported by the neural network represents the λ-message from the camera to the marketplace node in Figure 1. For uncertainty-aware processing using the more general subjective Bayesian network, the output of the Evidential Deep Learning network[10] provides the subjective opinion for the λ-message.

This paper accommodates the Evidential Deep Learning network a bit differently. It is assumed that the camera will be placed in one of two states based upon the marketplace activity. A camera observation does not provide absolute knowledge of the camera state value. Rather the observation is the subjective opinion output of the evidential neural network, and this becomes the λ-opinion, i.e., ω_{λ_X}, for the camera node.

4. APPROACHES TO EXPLAINING BAYESIAN MODELS

4.1 Background

Boerlage[11] was the first to formally introduce the concept of link strength for Bayesian Networks.[9] Given any pair of nodes—adjacent or not—their *connection strength* measures the strength of the conditional dependencies between them taking taking any possible path between them into account. In contrast, *link strength*, or *arc weight*, is defined for a specific edge and measures the strength of conditional dependencies only along that single edge.

In addition to being useful for visualisation of Bayesian networks[9] and for debugging, link strength measures have been applied to derive approximate inference to determine which variables can be neglected in the approximation.[15] Furthermore, it can also be used for structure learning and causality discovery.[16] For example, using constraint-based structure learning algorithms to learn a system's structure from data often yields large sets of Markov-equivalent DAGs. Using link strength can help researchers to narrow down the alternatives by eliminating those networks with only minor differences (i.e. those where arrows are reversed only for edges with very small link strength), thus yielding a more manageable number of major causal hypotheses to consider.[9]

The strength of the link from a random variable X to a random variable Y can be defined as the *mutual information*[9] of X, Y conditioned on the set $\mathbf{Z}$ of all other parents of Y:

$$\begin{aligned} LS^{\text{true}}(X \to Y) &= MI(X, Y|\mathbf{Z}) \\ &= \sum_{x,z} \sum_{y} P(y|x,z) \log_2 \tfrac{P(y|x,z)}{P(y|z)} \end{aligned} \tag{31}$$

4.2 Link Strength for SBN

We initially looked at approaches for estimating the edge strength between nodes in a network, i.e. the extent to which changes in the probability distribution at one node affect the probability distribution at a neighbouring node. This is a static measure that does not depend on any observed evidence, providing a global explanation that in effect summarises the conditional probability distributions in the network. We adapted (31) (hencefort $LS^{\text{true}}_{\text{SL}}(X \to Y)$) to be used with subjective opinions rather than standard probabilities by sampling from the beta distributions (which represent the uncertain opinions in our binary SLBNs) at the nodes and calculating an empirical link strength distribution. In this way the uncertainty in the link strength measure can be quantified as the spread/variance of the resulting link strength distribution.

4.3 Consistency Likelihood Edge Strength

Analysis in Section 5 of the various observations in the scenario reveals that the low probability, or novel, observations lead to high uncertainty for the inferred variables. This is because there is less relevant historical data to determine messages that propagate via the inferencing. On the other hand, when no observations are available, the uncertainty of the inferred variables are commensurate to the number of historical instantiations to determine the conditional opinions. This is because all of that data is relevant for the inference. Similarly, for common, or high probability, observations, the propagated messages should not change relative to the prior (or observed case) as the prior anticipates observations. We expect that the unexpected leads to high uncertainty, and it is the unexpected that cause the greatest changes in the messages during inference when a subset of variables are observed relative to the messages representing the case of no observed variable values.

Thus, we also developed a method for identifying "important" edges given observations for some of the nodes. By important, we mean the edges that created the biggest changes in opinions at connected nodes, and that had the greatest influence in changes in uncertainty, compared to the baseline case with no observations. This allows us to trace back from a particular query node to identify which opinions and edges were most influential in changing the opinion and uncertainty at that query node. For this we derived a formula for a likelihood ratio to compare a node's opinion before and after observing evidence. The value of the likelihood ratio formula indicates whether the prior and updated opinions (given evidence) are significantly different, similar, or if the difference between them is uncertain.

We implemented two different approaches to using this likelihood ratio method to measure edge importance. The first approach compares the opinions due to the observed evidence to the updated opinion given the baseline message (with no evidence) from each edge, one edge at at time. Specifically, the edge strength for $Z \to X$ is

$$ES^1(Z \to X) = \Lambda(\omega_X^{(e)-Z}, \omega_X^{(e)}) \tag{32}$$

where

$$\omega_X^{(e)-Z} = \left(\bigodot_{i=1}^{n} \omega_{\lambda_{Y_i,X}}^{(e)-Z} \right) \odot \left(\omega_{X|U_1 \ldots U_m} \circledcirc \omega_{\pi_{U_1,X}}^{(e)-Z} \ldots \omega_{\pi_{U_m,X}}^{(e)-Z} \right)$$

and the message opinions $\omega^{(e)-Z}$ are influenced by the observed evidence unless the sending variable is Z, i.e.,

$$\omega^{(e)-Z}_{\lambda_{Y_i,X}} = \begin{cases} \omega_{\lambda_{Y_i,X}} & \text{if } Z = Y_i \\ \omega^{(e)}_{\lambda_{Y_i,X}} & \text{otherwise} \end{cases} \quad \text{and} \quad \omega^{(e)-Z}_{\pi_{U_i,o_i}} = \begin{cases} \omega_{\pi_{U_i,o_i}} & \text{if } Z = U_i \\ \omega^{(e)}_{\pi_{U_i,o_i}} & \text{otherwise} \end{cases}$$

Note that ω and $\omega^{(e)}$ represent the baseline and evidence influenced opinions, respectively. This measure provides an indication of how much the change in each connected node affects the node under consideration independently of its other edges.

The second approach attempts to incorporate combined effects of edges by instead comparing the baseline opinion at each node (with no evidence) to the updated opinion given the evidence influenced message from each edge, one edge at a time. Specifically, the edge strength for $Z \to X$ is

$$ES^2(Z \to X) = \Lambda(\omega_X^{Z+(e)}, \omega_X), \tag{33}$$

where

$$\omega_X^{Z+(e)} = \left(\bigodot_{i=1}^{n} \omega^{Z+(e)}_{\lambda_{Y_i,X}} \right) \odot \left(\omega_{X|U_1 \ldots U_m} \circledcirc \omega^{Z+(e)}_{\pi_{U_1,X}} \ldots \omega^{Z+(e)}_{\pi_{U_m,X}} \right),$$

and the message opinions $\omega^{Z+(e)}$ are baseline opinions unless the sending variable is Z, i.e.,

$$\omega^{Z+(e)}_{\lambda_{Y_i,X}} = \begin{cases} \omega^{(e)}_{\lambda_{Y_i,X}} & \text{if } Z = Y_i \\ \omega_{\lambda_{Y_i,X}} & \text{otherwise} \end{cases} \quad \text{and} \quad \omega^{Z+(e)}_{\pi_{U_i,o_i}} = \begin{cases} \omega^{(e)}_{\pi_{U_i,o_i}} & \text{if } Z = U_i \\ \omega_{\pi_{U_i,o_i}} & \text{otherwise} \end{cases}$$

These two approaches often produce similar results but can differ depending on the conditional probability tables at the nodes and the configuration of the evidence nodes. We are currently conducting experiments to compare these methods to the naïve approach of removing single pieces of evidence one node at a time and recomputing opinions for the entire network, which for large networks is much more computationally expensive than our proposed methods.

5. EMPIRICAL ANALYSIS

As discussed in Section 2, the actual conditional probabilities of the Bayesian network (cf. Figure 1) are not known to the GoodGuys. Instead, they learn a subjective Bayesian network using human intelligence that obtained N_{ins}=100 past daily instantiations of the dispositions of the two factions, along with reports of route conditions and civic center attendance. The ground truth conditional probabilities for the Bayesian network is given in Table 1.* The table indicates that the Capulets usually have a favorable disposition towards the Goodguys, in contrast to the Montagues. Whenever either faction has a negative disposition, attendance at the corresponding civic center is higher than normal. It is also more likely the faction will disrupt any Goodguys traveling through their territory. Violent activities in the marketplace are only likely to occur when both factions have negative dispositions, and only then is it likely the factions would disrupt GoodGuys traveling along Route B. The commander's knowledge of all these trends is based on the Subjective Bayesian network obtained from the N_{ins}=100 instantiations.

Table 2 provides the inferred opinions of the route conditions from the subjective Bayesian network for each of the three routes for various observations. In the absence of any observations (Scenario 1, Table 2), the uncertainty of all there opinions is relatively low and reflective of 100 effective pieces of evidence. The opinions for routes A and B are similar, and these beliefs are flipped for the Route C opinion. These opinions are consistent with general trends from the ground truth conditional probabilities.

When the camera is detecting normal activity at the marketplace, and the attendance at the civic centers in Capulet and Montague sections are high and normal, respectively Scenario 2 and 3 of Table 2, the route conditions opinions demonstrate low uncertainty. Since these observations are likely to occur, the uncertainty remains small, and observations increases the

*In this example, the variables are binary simply because the current subjective belief propagation method works only for the binary case. We are currently extending the propagation method to accommodate more general multinomial variables.

Variable	States	Conditional Probabilities
Capulet Disposition	$\{\texttt{pos}, \texttt{neg}\}$	$P(\text{CD} = \texttt{pos}) = 0.9$
Montague Disposition	$\{\texttt{pos}, \texttt{neg}\}$	$P(\text{MD} = \texttt{pos}) = 0.1$
Capulet Civic Center Attendance	$\{\texttt{norm}, \texttt{high}\}$	$P(\text{CCA} = \texttt{norm} \mid \text{CD} = \texttt{pos}) = 0.8$ $P(\text{CCA} = \texttt{norm} \mid \text{CD} = \texttt{neg}) = 0.1$
Montague Civic Center Attendance	$\{\texttt{norm}, \texttt{high}\}$	$P(\text{MCA} = \texttt{norm} \mid \text{MD} = \texttt{pos}) = 0.8$ $P(\text{MCA} = \texttt{norm} \mid \text{MD} = \texttt{neg}) = 0.1$
Marketplace Activity	$\{\texttt{norm}, \texttt{violent}\}$	$P(\text{Ma} = \texttt{norm} \mid \text{CD} = \texttt{pos} \wedge \text{MD} = \texttt{pos}) = 0.99$ $P(\text{Ma} = \texttt{norm} \mid \neg(\text{CD} = \texttt{neg} \wedge \text{MD} = \texttt{neg}) = 0.01)$
Camera	$\{\texttt{norm}, \texttt{violent}\}$	$P(\text{Cam} = \texttt{norm} \mid \text{Ma} = \texttt{norm}) = 0.95$ $P(\text{Cam} = \texttt{norm} \mid \text{Ma} = \texttt{violent}) = 0.01$
Route A Condition	$\{\texttt{safe}, \texttt{danger}\}$	$P(\text{RA} = \texttt{safe} \mid \text{CD} = \texttt{pos}) = 0.9$ $P(\text{RA} = \texttt{safe} \mid \text{CD} = \texttt{neg}) = 0.1$
Route B Condition	$\{\texttt{safe}, \texttt{danger}\}$	$P(\text{RB} = \texttt{safe} \mid \text{Ma} = \texttt{norm}) = 0.9$ $P(\text{RB} = \texttt{safe} \mid \text{Ma} = \texttt{violent}) = 0.1$
Route C Condition	$\{\texttt{safe}, \texttt{danger}\}$	$P(\text{RC} = \texttt{safe} \mid \text{MD} = \texttt{pos}) = 0.9$ $P(\text{RC} = \texttt{safe} \mid \text{MD} = \texttt{neg}) = 0.1$

Table 1: Ground truth conditional probabilities for the Bayesian network in Figure 1.

beliefs that routes A and B are safe and Route C is dangerous. With these observations, the commander can choose Route A or B with limited armor support.

When the attendance at the Montague civic centers goes normal, then uncertainty about Route C increases while the opinions of routes-A and B are unaffected. This is due to the fact that it is unlikely for the Montagues to have a positive disposition towards the GoodGuys, which reduces the amount of effective training samples to determine the condition or Route C. Again the commander can choose Route A or B with limited armor support.

Then when the camera has high confidence of violence at the marketplace (Scenario 4, Table 2), the uncertainty for the opinions of routes B and C increases drastically. This is due the fact that it is extremely unlikely for violence to occur in the marketplace when attendance at both civic centers are at normal levels. There are very few training instantiations to help determine these opinions. On the other hand, the uncertainty of the opinion for Route A has not increased as much because the Capulet civic center activity in not as unexpected as the two other observations. Nevertheless, the uncertainty for Route A has doubled, and the commander would need to pick it; but perhaps requesting heavier armor support.

Finally, the neural network reports that it is uncertain about the activity class (Scenario 5, Table 2). This can lead the commander to direct his team to analyze the video footage. The team determines that the people in the marketplace have synchronized in a dance routine that the neural network was never trained to recognize. The commander digs further with some experts to learn that the dance is actually an ancient ritual that two tribes in the region perform to synchronize the mind and body when the tribes are preparing to attack a common exogenous enemy. In light of this additional context, the commander decides to provide heavy armed support on both the ground and air to escort the supply truck.

The scenario illustrates how understanding uncertainty is important for the commander to assess risk. The uncertainty can arise because the observations are due to rare events for which sparse training means few exemplars to establish opinions for the states of the decision variables. The uncertainty can also arise when the neural networks cannot interpret the input data. This case can represent a ‘black swan’ event that the neural network was never trained to understand. The synchronized dance is an example of an unknown unknown that requires human investigation.

Scenarios	Route A			Route B			Route C		
	$b_{\texttt{safe}}$	$b_{\texttt{dang}}$	u	$b_{\texttt{safe}}$	$b_{\texttt{dang}}$	u	$b_{\texttt{safe}}$	$b_{\texttt{dang}}$	u
Scenario 1 No observations	0.76	0.22	0.02	0.82	0.16	0.02	0.22	0.76	0.02
Scenario 2 CCA = `norm` MCA = `high` $\omega_{\text{Ma}}^{\text{Cam}} = (0.95, 0, 0.05)$	0.85	0.13	0.02	0.92	0.06	0.02	0.17	0.81	0.02
Scenario 3 CCA = `norm` MCA = `norm` $\omega_{\text{Ma}}^{\text{Cam}} = (0.95, 0, 0.05)$	0.84	0.14	0.02	0.92	0.06	0.02	0.51	0.39	0.10
Scenario 4 CCA = `norm` MCA = `norm` $\omega_{\text{RA}}^{\text{Ma}} = (0, 0.95, 0.05)$	0.76	0.19	0.05	0.51	0.30	0.19	0.54	0.31	0.15
Scenario 5 CCA = `norm` MCA = `norm` $\omega_{\text{Ma}}^{\text{Cam}} = (0, 0, 1)$	0.83	0.15	0.02	0.84	0.09	0.07	0.52	0.38	0.10

Table 2: Opinions for the route conditions obtained via subjective belief propagation in light of various observations. Note that camera likelihood opinion $\omega_{\text{RA}}^{\text{Cam}} = (b_{\texttt{norm}}, b_{\texttt{violence}}, u)$ is the backpropagated message to the marketplace activity variable from the camera data.

5.1 Link Strength

Figure 2 and Table 3 (in the Appendix) report the results of computing $LS_{\text{SL}}^{\text{true}}$ for each of the links of the SLBN depicted in Figure 1, instantiated with conditional probabilities as presented in Table 1. From the analysis—that is based only on the static component of the network—it is evident that the link between Ma and Cam is the strongest, followed by Ma → RB, and then MD → RC.

5.2 Consistency Likelihood Edge Strength

While $LS_{\text{SL}}^{\text{true}}$ considers only the static conditional probability tables, the two approaches fo consistency likelihood edge strength (Section 4.3) do take into consideration evidence injected into the network. Figure 3 provides a graphical representation of ES^1 for each edge of our SBN in the four scenarios—cf. Table 2—where evidences are injected in it: actual values are present in Table 4 in the appendix. Similarly, Figure 4 (and Table 5 in the appendix) summarises the results for ES^2.

Although the two approaches naturally differ, they provide some consistent guidance. As per Scenario 2 (cf. Table 2), Figures 3a and 4a show how the link Ma → Cam is among the most important here. Even more marked the similarities between Figure 3b and Figure 4b; Figure 3c and Figure 4c; and Figure 3d and Figure 4d. In there, it is evident how the links MD → MCA and MD → MCA play a significant role. However, it is worth noticing that the link Ma → Cam is acknowledged to play a significant role in Scenarios 2,3, and 4 when using ES^2—perhaps a reflection of the confidence and the conclusiveness of the conditional opinions as highlighted in Table 2—while it is not very significant as per Scenario 3 when using ES^1 (cf. Figure 4b). This difference will be investigated in more details as part of future work.

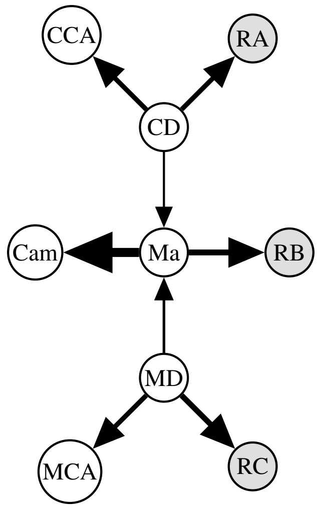

Figure 2: Graphical representation of $LS_{\mathrm{SL}}^{\mathrm{true}}$ for each edge in the SLBN depicted in Figure 1. The thicker the line, the stronger is the link.

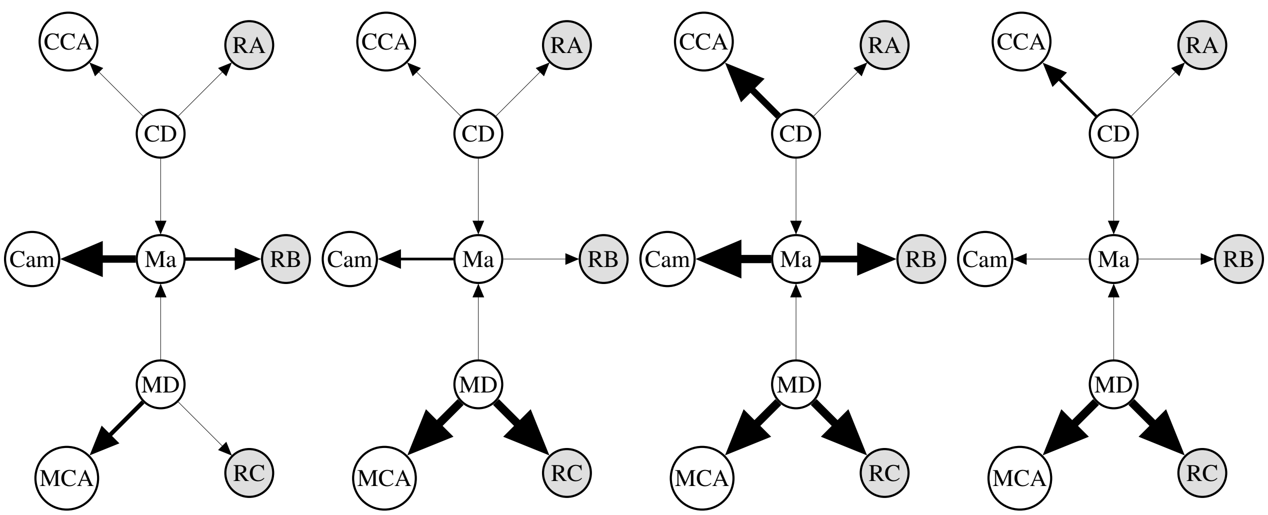

(a) Scenario2 (b) Scenario3 (c) Scenario4 (d) Scenario5

Figure 3: Graphical representation of ES^1 for each edge in the SLBN depicted in Figure 1 for each of the scenarios described in Table 2.

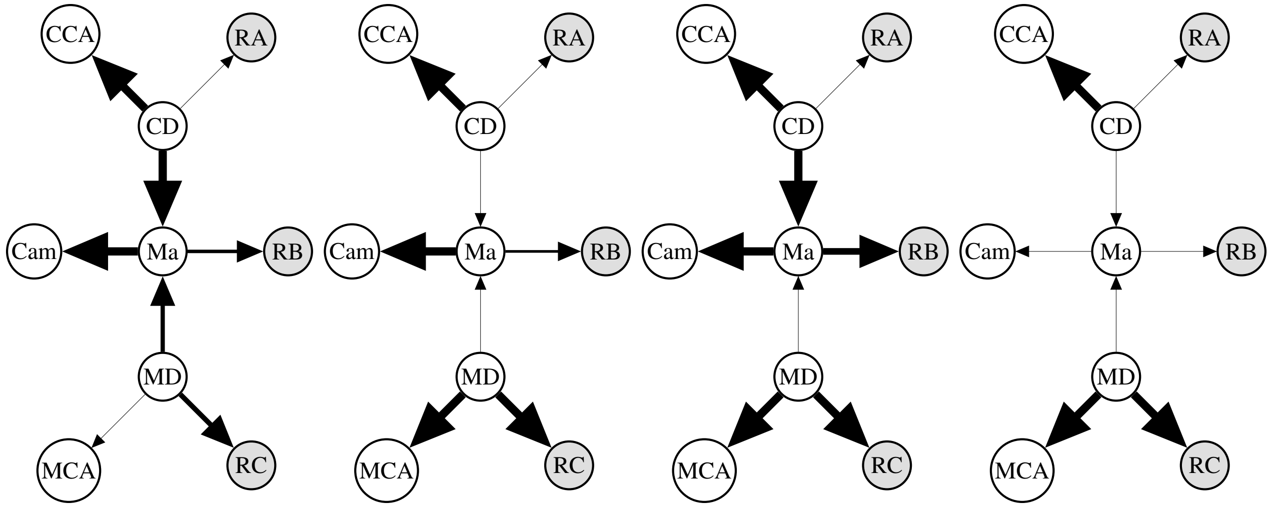

(a) Scenario2 (b) Scenario3 (c) Scenario4 (d) Scenario5

Figure 4: Graphical representation of ES^2 for each edge in the SLBN depicted in Figure 1 for each of the scenarios described in Table 2.

6. CONCLUSIONS

In this paper we argue that situational understanding needs reasoning both under and about uncertainty capabilities. In particular, we considered the case of Subjective Bayesian Network, and we provided a case-study analysis of techniques for strength analysis in SBN (Section 4). We derived $LS_{\mathrm{SL}}^{\mathrm{true}}$, an approach based on mutual information between two random variables originally proposed in the context of Bayesian networks,[9] and we compared it with two native methods exploiting the uniqueness of the SBN approach, ES^1 and ES^2. We then discussed an articulated empirical analysis (Section 5) that shows strengths and weaknesses of our three methods for analysing the strengths of the causal links in the network.

In summary, in our case study, $LS_{\mathrm{SL}}^{\mathrm{true}}$ does seem to highlight the edges that connect the query variables of high uncertainty to the *unusual* observations. It seems inappropriate for such a task as it only determines whether or not the correlation between variables are strong or not. It does not consider the *novelty* of inference messages when some variable values are observed.

Considering ES^1 and ES^2, it appears that there is an inverse correlation between the degree of aleatoric uncertainty of the evidence and the edge strength. However, this is not straightforward to see in ES^1. We are therefore committed to future investigation to shed more light on such correlations.

Acknowledgement

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

[1] Dostal, B. C., "Enhancing situational understanding through employment of unmanned aerial vehicle," *Army Transformation Taking Shape: Interim Brigade Combat Team Newsletter* **01-18** (2007).

[2] "Understandig: Joint Doctrine Publication 04 (JDP 04)." Ministry of Defence, UK (2016).

[3] Cerutti, F., Alzantot, M., Xing, T., Harborne, D., Bakdash, J. Z., Braines, D., Chakraborty, S., Kaplan, L. M., Kimmig, A., Preece, A. D., Raghavendra, R., Sensoy, M., and Srivastava, M. B., "Learning and reasoning in complex coalition information environments: A critical analysis," in [*21st International Conference on Information Fusion, FUSION 2018, Cambridge, UK, July 10-13, 2018*], 1–8 (2018).

[4] Hinton, G. E., Osindero, S., and Teh, Y.-W., "A fast learning algorithm for deep belief nets," *Neural computation* **18**(7), 1527–1554 (2006).
[5] Krizhevsky, A., Sutskever, I., and Hinton, G. E., "ImageNet classification with deep convolutional neural networks," in [*Advances in neural information processing systems*], 1097–1105 (2012).
[6] LeCun, Y., Bengio, Y., and Hinton, G., "Deep learning," *Nature* **521**(7553), 436–444 (2015).
[7] M.Kaplan, L. and Ivanovska, M., "Efficient belief propagation in second-order Bayesian networks for singly-connected graphs," *International Journal of Approximate Reasoning* **93**, 132–152 (Feb. 2018).
[8] Kaplan, L. M., Cerutti, F., Sensoy, M., Preece, A. D., and Sullivan, P., "Uncertainty aware AI ML: why and how," *CoRR* **abs/1809.07882** (2018).
[9] Ebert-Uphoff, I., "Tutorial on How to Measure Link Strengths in Discrete Bayesian Networks," tech. rep. (Sept. 2009).
[10] Sensoy, M., Kandemir, M., and Kaplan, L., "Evidential deep learning to quantify classification uncertainty." `https://arxiv.org/abs/1806.01768` (2018).
[11] Kotz, S., Johnson, N. L., and Balakrishnan, N., [*Continuous multivariate distributions.*], Wiley (2000).
[12] Jøsang, A., [*Subjective Logic: A Formalism for Reasoning Under Uncertainty*], Springer (2016).
[13] Kaplan, L. and Ivanovska, M., "Efficient belief propagation in second-order Bayesian networks for singly-connected graphs," *International Journal of Approximate Reasoning* **93**, 132–152 (2018).
[14] Pearl, J., "Fusion, propagation, and structuring in belief networks," *Artificial Intelligence* **29**(3), 241–288 (1986).
[15] Jitnah, N. and Nicholson, A., "treenets: A framework for anytime evaluation of belief networks," in [*Qualitative and Quantitative Practical Reasoning*], 350–364, Springer (1997).
[16] Ebert-Uphoff, I., Deng, Y., Ebert-Uphoff, I., and Deng, Y., "Causal Discovery for Climate Research Using Graphical Models," *Journal of Climate* **25**, 5648–5665 (sep 2012).

APPENDIX A. TABLE OF RESULTS

	Scenario 1	
	Mean	Std
CD → CCA	0.18	0.08
CD → RA	0.18	0.08
CD → Ma	0.08	0.06
Ma → Cam	0.35	0.1
Ma → RB	0.22	0.08
MD → Ma	0.1	0.08
MD → MCA	0.18	0.08
MD → RC	0.21	0.09

Table 3: $LS_{\text{SL}}^{\text{true}}$ for each edge in the SLBN depicted in Figure 1 for each of the scenarios illustrated in Table 2.

Message update	**Scenario 2**	**Scenario 3**	**Scenario 4**	**Scenario 5**
CCA → CD	5.0	2.0	0.7	1.3
CD → RA	2.3	2.7	5.4	3.6
CD → Ma	20.9	21.6	1.6	3.0
Ma → CD	6.6	13.1	2.3	12.8
Cam → Ma	0.6	1.4	0.3	4.3
Ma → RB	1.1	1.2	0.7	5.0
Ma → MD	19.2	2.3	2.0	2.3
MD → Ma	49.9	33.7	1.8	4.3
MCA → MD	1.0	0.0	0.1	0.0
MD → RC	4.2	0.1	0.1	0.1

Table 4: ES^1 for each vertex pair in the SLBN depicted in Figure 1, for each of the scenarios illustrated in Table 2

Message update	**Scenario 2**	**Scenario 3**	**Scenario 4**	**Scenario 5**
CCA → CD	0.5	0.5	0.5	0.5
CD → RA	2.3	2.7	5.4	3.5
CD → Ma	1.6	1.6	1.6	1.6
Ma → CD	0.3	3.7	0.3	6.5
Cam → Ma	0.0	0.0	0.0	4.2
Ma → RB	1.1	1.2	0.7	5.0
Ma → MD	9.6	9.6	5.7	9.5
MD → Ma	8.5	6.6	6.6	6.6
MD → RC	4.2	0.1	0.1	0.1
MCA → MD	0.9	0.0	0.0	0.0

Table 5: ES^2 for each vertex pair in the SLBN depicted in Figure 1, for each of the scenarios illustrated in Table 2