Meta-Reasoning about Decisions in Autonomous Semi-Intelligent Systems

Mats Danielson
Dept. of Computer and Systems Sciences
Stockholm University
PO Box 7003
SE-164 07 Kista, SWEDEN
+468 161 540
mats.danielson@su.se

Love Ekenberg
International Institute of
Applied Systems Analysis
Schlossplatz 1
A-2361 Laxenburg, Austria
+468 161 679
ekenberg@iiasa.ac.at

ABSTRACT
For intelligent systems to become autonomous in any real sense, they need an ability to make decisions on situations that were not entirely conceived of at compile-time. Machine learning algorithms are excellent in mimicking the behaviour of some gold standard role model, and this can include decision making by the role model. But once out of familiar contexts, the decision making becomes harder and needs an element of more independent probabilistic reasoning and decision making. This paper presents such a method based on a belief mass interpretation of the decision information, where the components are imprecise and thus uncertain by means of intervals.

CCS Concepts
• **CCS → Applied computing → Operations research → Decision analysis → Decision-making**

• **CCS → Applied computing → Operations research → Decision analysis → Multi-criteria optimization**

Keywords
Autonomous intelligent systems; decision analysis; meta-reasoning; imprecise probabilities.

1. INTRODUCTION
In the coming digital machine age, more and more decisions will be made by non-human entities. This translates, in the vast majority of cases, to software running on digital silicon hardware. We are witnessing a surge in the field of artificial intelligence (AI), a field with concepts of which there are many definitions. The early (pre-1980) definitions of strong and weak AI notwithstanding, AI has over the years meant and still is taken to mean vastly different approaches and phenomena. For some, it mostly encompasses machine learning algorithms and training, sometimes even focusing on deep learning above all. For others, it contains the majority of digitalisation witnessed in society, ranging from all administrative data processing systems, over all internet activities, to Internet of things. Trying to cover the middle ground, we will focus on intelligent autonomous or semi-autonomous systems running some kind of information gathering software. It could be sensors connected to different kinds of inference machines, but it could also be other means of inferring knowledge from data. One problem with some approaches is that they mix inference with decisions. If a system lets a learning algorithm train on an adequate and sufficiently large set of data, the outcome will be a (hopefully) good categorisation of the situation at hand. But some approaches let the system train not on categories of situations but rather directly on decisions. This, unfortunately, leads to less flexible and less transparent system properties since the detection, categorisation, and decision steps are all tied together and cannot be used or improved on separately. In this paper, we suggest a decision-making mechanism that is open and transparent and is well-suited for processing the output from detection and categorisation algorithms that deal with imprecise and probabilistic data, even taking different criteria (different categorisations of the same situations from different perspectives) into account. Output from Bayes' nets, fuzzy inference systems, and traditional machine learning all fall into this set of feasible input, the latter at least in the cases where the networked learned can express its certainty of the inference reached by some measure.

We will in this paper discuss the representation and evaluation of second-order information in order to allow autonomous entities to make decisions based on incomplete input data in the form of imprecise utilities, probabilities, and criteria weights. The approach avoids the introduction of new concepts into the decision models, such as set membership functions or similar formalisms, and instead use second-order distributions of belief in the basic utilities, probabilities, and criteria weights which then allows for better and more transparent discrimination between the resulting values of the decision alternatives. The ability to use belief distributions over probabilities, values, and criteria weights enhances the transparency of the results since no new concepts have to be introduced in the evaluation of the model.

2. BACKGROUND
In the research community, there have been many suggestions as to how to handle the very strong requirements for decision-makers, whether machine or human, to provide precise information. Some main categories of approaches to remedy the precision problem are based on capacities, sets of probability measures, upper and lower probabilities, interval probabilities (and sometimes utilities), evidence and possibility theories, as well as fuzzy measures (see, for example, [1-5] for different suggestions). The latter category seems to be used only to a limited extent in real-life decision analyses since it usually requires a significant mathematical background on the part of the decision-maker. Another reason is that computational complexity can be problematic if the fuzzy aggregation mechanisms are not significantly simplified.

A main issue with the above approaches is that they provide very little assistance when the results overlap, as is usually the case in real-life decision problems. This issue will be addressed in the paper. In Figure 1, the uncommon case of no overlap is illustrated. In the figure and the following figures, the belief distribution in the expected value is shown, i.e. with expected value on the x-axis and belief density on the y-axis. Since in Figure 1 all the belief in Action-1 is at higher expected values than that in Action-2, i.e., there is no chance that the expected value of Action-2 would be higher than that of Action-1, provided we accept the principle of maximizing the expected value. Thus, there is no doubt that the semi-intelligent system (SIS) should select Action-1.

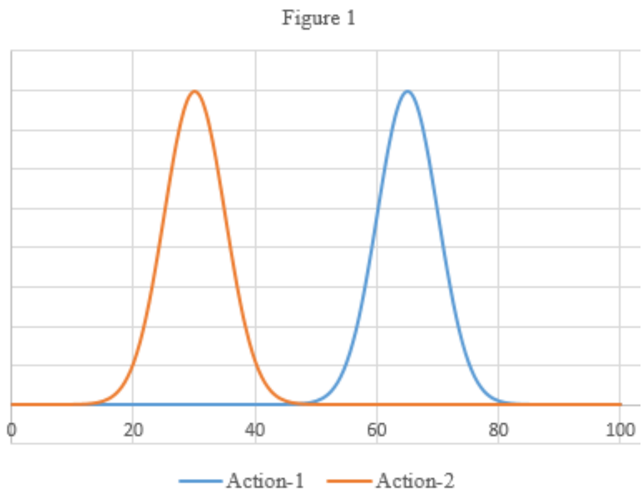

Figure 1. No overlap in expected value

In Figure 2, the more common case of some overlap is illustrated. Since the belief in Action-1 is still significantly higher than that in Action-2, there is no reasonable doubt that the SIS should, also in this case, select Action-1.

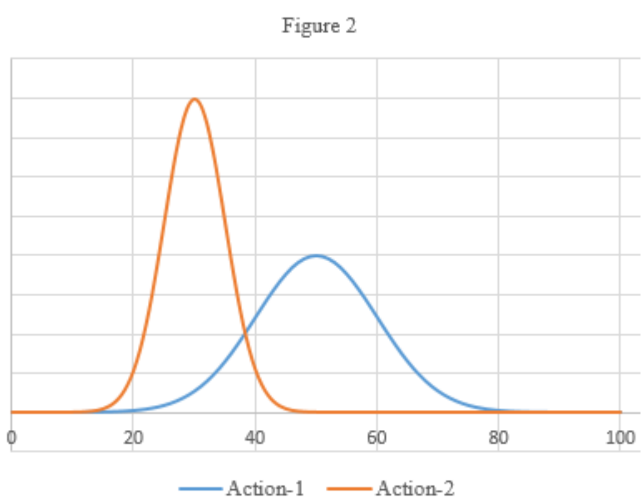

Figure 2. Minor overlap in expected value

In Figure 3, on the other hand, the common case of large overlap is illustrated. Since now the belief in Action-1 and Action-2 overlap a lot, there is no clear-cut decision available as to whether the SIS should select Action-1or Action-2 in this case. In the following sections, we will suggest a computational model based on higher-order belief for discriminating between actions in all of the cases above, especially focussing on the situation in Figure 3. We begin with the representation of knowledge before moving on to the evaluation of the representational model.

[1] With *decisions* we are in this paper referring to larger, strategic choice problems rather than mainly reactive responses to stimuli.

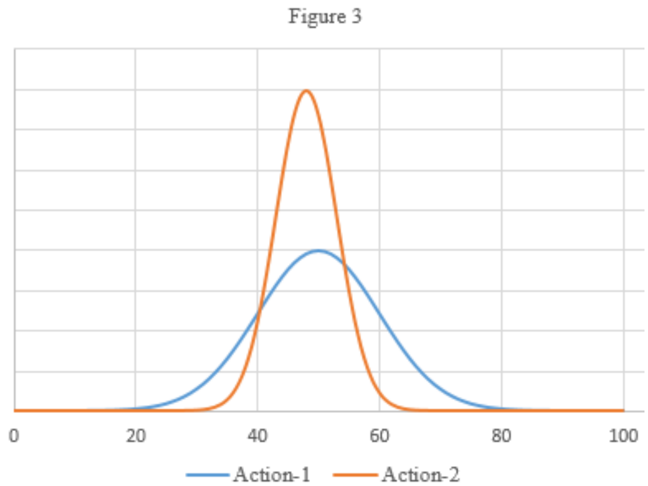

Figure 3. Major overlap in expected value

3. HIGHER ORDER BELIEF

Let us assume that we have an SIS, i.e. a system that is capable of handling situations that were not entirely conceived of at compile-time. In order to be semi-intelligent under any reasonable interpretation of the terminology, it should be able to make decisions[1] on its own. Further assume that it has a set of alternative actions at its disposal, from which the SIS is to choose one. The alternatives are being considered using information from a set of criteria. Thus, each alternative is assessed under each criterion. The criteria are then combined together using a weighting scheme under the traditional assumption of additivity, but to simplify the presentation in this paper we will only discuss the single-criterion case in which the expected value rule is to maximize the expected value of an alternative

$$E(\mathrm{A}_i) = \sum_k p_{ik} \cdot v_{ik}$$

where all the products $p_{ik} \cdot v_{ik}$ of an alternative A_i are added to form the sum.

The expected value of each alternative under each criterion is represented by a decision tree. The components of such a decision tree are a root node, a set of probability nodes and consequence nodes. The probability and consequence nodes are in a standard model assigned unique numerical probability and value distributions. When an alternative A_i is chosen as the preferred action, there is a probability p_{ij} that an event occurs that leads either to another subsequent event or to a consequence with a value v_{ijk}.

The SIS will in most cases, when operating in real-life contexts, not be able to acquire precise or complete information with respect to the probabilities and utilities. There will be information available but with uncertainty involved, and the idea is to use this uncertainty to our advantage by expressing it in a machine computable way. In [6], we describe a formalism for handling second-order belief by using a modified normal distribution of belief over the input data. Note that this is the distribution of belief in numerical assessments (i.e. second-order information) by the SIS.

Let us further assume that the SIS has been able to establish intervals within which the particular decision variables (probabilities and utilities/values) lie. The beliefs in the numbers assigned to those variables are not uniformly distributed. Rather, the belief in outer boundary numbers is modelled to be less than in more central numbers [7-8], and since it is plausible to presume that the SIS has a higher belief in more central parts of its derived intervals for probabilities and utilities/value, the idea is to represent belief as distributions over those probabilities and values.

The SIS should employ different distributions for probabilities and values because of the normalisation constraints for probabilities. For values, which is the simpler case, ordinary triangular distributions will do. They constitute a good-enough centre-weighted representation of the SIS's belief.

Definition: A unit cube is all tuples $(x_1, \ldots, x_n)$ in $[0,1]^n$ and a second-order distribution over that cube B is a positive distribution F defined on B such that

$$\int_B F(x)\, dV_B(x) = 1,$$

where V_B is the n-dimensional Lebesgue measure on B.

For probabilities, a natural candidate is the Dirichlet distribution. The properties of the Dirichlet distribution, being a parameterised family of continuous multivariate probability distributions, makes it suitable for this purpose. The probability density function of the Dirichlet distribution is defined as

$$f_{dir}(p,\alpha) = \frac{\Gamma(\sum_{i=1}^{k} \alpha_i)}{\prod_{i=1}^{k} \Gamma(\alpha_i)} p_1^{\alpha_i - 1} p_2^{\alpha_2 - 1} \ldots p_k^{\alpha_k - 1}$$

on a set $\{p=(p_1,\ldots p_k) | p_1,\ldots,p_k \geq 0, \Sigma p_i = 1\}$, where $(\alpha_1,\ldots,\alpha_k)$ is a parameter vector in which each $\alpha_i > 0$ and $\Gamma(\alpha_i)$ is the Gamma function. The Dirichlet distribution is a multivariate generalisation of the beta distribution, and the marginal distributions of Dirichlet are indeed beta distributions.

The SIS will use a slightly different form due to the varying sizes of the intervals over which the belief will be distributed. This is the bounded Dirichlet distribution which is defined over a particular range instead of the standard interval [0,1]. Bounded beta distributions are then derived from this, yielding four-parameter beta distributions. Thus, we use a probability belief distribution by employing a bounded Dirichlet distribution $f_3(a_i, c_i, b_i)$ where c_i is the estimated most likely probability and where a_i and b_i are the boundaries for the support of the distribution ($a_i < c_i < b_i$) (cf. [9]).

4. AGGREGATION PRINCIPLES

To evaluate the decision situation means to determine the resulting distribution over the expected utility/value. Thus, there are only two cases to cover – multiplication and addition.

Definition: Let G be a distribution over two cubes A and B, and assume that G has positive support for the feasible distributions at level i in a decision tree, and for the feasible probability distributions of the children of a node x_{ij}. Also assume that $f(x)$ and $g(y)$ are marginal distributions of $G(z)$ on A and B, respectively. Then the cumulative multiplied distribution of the two belief distributions is

$$\mathrm{H}(z) = \iint_{\Gamma_z} f(x)g(y)dxdy = \int_0^1 \int_0^{z/x} f(x)g(y)dxdy = \int_z^1 f(x)G(z/x)dx$$

where G is a primitive function of g, $\Gamma_z = \{(x,y) \mid x \cdot y \leq z\}$, and $0 \leq z \leq 1$.

Let $h(z)$ be the corresponding density function. Then

$$h(z) = \frac{d}{dz}\int_z^1 f(x)G(z/x)dx = \int_z^1 \frac{f(x)g(z/x)}{x}dx.$$

Addition is handled by the standard convolution of two densities, restricted to A and B and the distribution h on a sum $z = x+y$, where we have the belief distributions $f(x)$ and $g(y)$ is consequently given by

$$h(z) = \frac{d}{dz}\int_0^z f(x)g(z-x)dx.$$

Employing this iteratively, the SIS will obtain the distribution over the expected value (or expected utility). See [10-11] for a more detailed account of this.

5. BELIEF DOMINANCE

In most cases, the resulting distributions of belief in the expected values will significantly overlap. (For cases where they do not overlap, any simple mechanism would be able to pick out the preferred choice.) If it were a human that was the decision-maker, that person could ponder over the resulting beliefs and, perhaps using a software tool with graphic presentation of the results, study and assess the resulting beliefs. But an SIS must be able to make those decisions in a stand-alone, automated way. To achieve this, we suggest using the concept of belief dominance. In the following, we will discuss the case of a single criterion, but it can without loss of generality be extended to the multi-criteria case.

Definition: Given a decision problem P with two possible actions A_i and A_j, let δ_{ij} denote the expression $\sum k\ p_{ik} \cdot v_{ik} - \sum k\ p_{jk} \cdot v_{jk}$ over all consequences in the consequence sets that make up the alternative courses of action A_i and A_j in P.

The index set pair captures the consequences within each of the alternatives that possess some desired property, in this case their value being at least as great as a given parameter.

Definition: Given a decision problem P and a real number d $\in$ [0,1], an *index set pair* $\langle K_i,K_j \rangle(d)$ in P is a pairing of two sets $K_i = \{k \mid v_{ik} \geq d\}$ and $K_j = \{k \mid v_{jk} \geq d\}$.

The parameter d varies and this represents a selection procedure for selecting the consequences within a pair of alternatives with the desired property.

Definition: Given a decision problem P and real numbers a,b,d $\in$ [0,1], $M_{ij}[a,b]$ is the set $\{\langle K_i,K_j \rangle(d) \mid d \in [a,b]\}$ in P.

Thus, $M_{ij}[a,b]$ is the set of all different index set pairs in the range [a,b], i.e. all the combinations of index sets that satisfy the condition. Now we can define belief dominance.

Definition: Given a decision problem P, a function f, and two parameters α and β, an alternative A_i *B-dominates* another alternative A_j in P **iff**

$\forall \langle K_i,K_j \rangle(d) \in M_{ij}[I]$

$$\sum_{k \in K_i} f(p_{ik}, v_{ik}, \alpha) - \sum_{k \in K_j} f(p_{jk}, v_{jk}, \beta) \geq 0,$$

and

$\exists \langle K_i,K_j \rangle(d) \in M_{ij}[I]$

$$\sum_{k \in K_i} f(p_{ik}, v_{ik}, \alpha) - \sum_{k \in K_j} f(p_{jk}, v_{jk}, \beta) > 0$$

From the definition of general belief dominance, we now derive operational concepts that an SIS could employ in separating alternatives with overlapping beliefs.

First-order belief dominance means that the belief dominance function used is a function of only probabilities, i.e. $f(p_{ik},v_{ik},\alpha)$ above is a function $g(p_{ik})$ and likewise, $f(p_{jk},v_{jk},\beta)$ is a function $g(p_{jk})$. Further, if the range for the index set pairs is [0,1], we arrive at first-order belief dominance.

Definition: Given a decision problem P, an alternative A_i *B_1-dominates* another alternative A_j in P **iff** Ai B-dominates A_j with I = [0,1] and $f(p_{ik},v_{ik},\alpha) \equiv g(p_{ik}) \equiv p_{ik}$ and likewise for p_{jk}.

Second-order belief dominance means that the belief dominance function used is a function of both probabilities and values i.e. $f(p_{ik},v_{ik},\alpha)$ above is a function $h(p_ik,v_{ik})$ and likewise $f(p_{jk},v_{jk},\beta)$ is a function $h(p_{jk},v_{jk})$. Also in this case, the range for the index set pairs is [0,1].

Definition: Given a decision problem P, an alternative A_i *B_2-dominates* another alternative A_j in P **iff** A_i B-dominates A_j with I = [0,1] and $f(p_{ik},v_{ik},\alpha) \equiv h(p_{ik},v_{ik}) \equiv p_{ik}\cdot v_{ik}$ and likewise for p_{jk}.

The ordinary expected value can also be seen as second-order belief dominance evaluated only by fully indexed pairs, i.e. pairs that contain all members of each alternative.

Definition: Given a decision problem P, an alternative A_i *B_E-dominates* another alternative A_j in P **iff** A_i B-dominates A_j with I = [0,0] and $f(p_{ik},v_{ik},\alpha) \equiv h(p_{ik},v_{ik}) \equiv p_{ik}\cdot v_{ik}$ and likewise for p_{jk}.

By inspection, B_E-dominance is found to be corresponding to an evaluation rule that applies a probability and value-based formula to the alternative in order to reach a single-numbered numerical verdict on which one is to prefer. This is the ordinary expected value since the only index set pair generated by the [0,0]-range is the pair of complete consequence sets. To sum up, an SIS now could have a more complete set of decision rules at its disposal than if employing only standard techniques of expected utility and non-overlapping intervals. When the SIS encounters overlap, it applies B_1-dominance to see if that results in a sufficient ranking. In that case, a strong ranking is obtained. Otherwise, the next step is to apply B_2-dominance to see if that results in a sufficient ranking, which then is a weak ranking but still better than trying to use the expected value which is the last resort if a decision must be made and all belief dominance techniques come up empty-handed.

6. EXAMPLE

Assume that an autonomous SIS agent entity has encountered the following imprecise and uncertain decision situation. It has to select between three actions, viz. Action-1, Action-2, and Action-3. For all three actions, its sensors and actuators have reported a lot of data that has been interpreted and aggregated into a decision situation in which each action is represented by an event tree containing uncertain interval probabilities and interval utilities/values. From this set of information, the agent computes interval-valued expected values which unfortunately overlap significantly. Thus, the agent entity is not able to choose between the alternative courses of action without further consideration. Assessing its belief in the probabilities and values, it uses the method in the preceding sections to arrive at distributions of belief over the actions' respective expected value intervals. It starts with comparing the belief of Action-1 and Action-2 by studying belief dominance. In Figure 4, the distribution of belief is shown and in Figure 5, the accumulated distribution of belief is shown. We can see that although the belief overlaps considerably, the belief in Action-1 is stronger since Action-1 B_1-dominates Action-2. This is seen by the fact that for each level of accumulated belief (y-axis), the sum of beliefs up to that point is for higher expected values for Action-1, i.e. its curve is always to the right, having higher expected value for any level of accumulated belief ranging from none (i.e. 0.00) to all (i.e. 1.00). Thus, Action-2 is found to be inferior to Action-1 and will not be considered by the SIS any further.

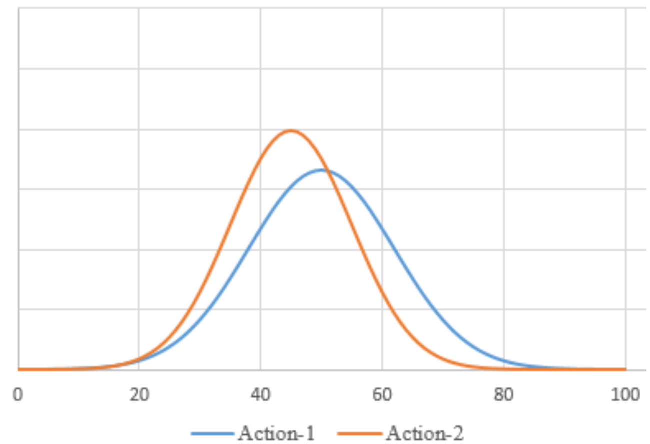

Figure 4. Comparing Action-1 and Action-2 (distribution of belief in expected values)

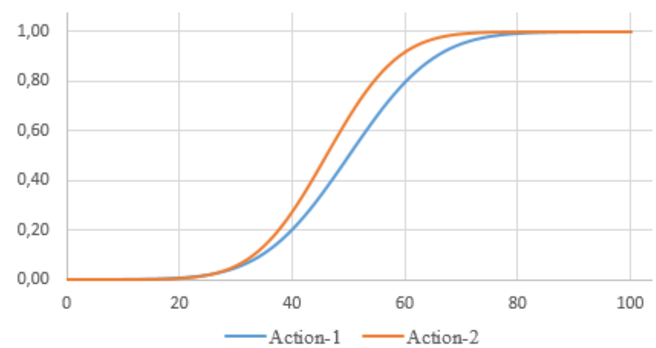

Figure 5. Comparing Action-1 and Action-2 (accumulated distribution of belief in expected values)

Next, the SIS agent entity continues with comparing the belief of Action-1 and Action-3 by studying belief dominance. In Figure 6, the distribution of belief is shown and in Figure 7, the accumulated distribution of belief is shown. We can see that in this case, the Action-1 does not B_1-dominate Action-3 since in the comparison in Figure 7, the accumulated curves cross. So the SIS has to move on to the weaker comparison tool of B_2-dominance. Comparing the two actions in this way, the belief in Action-1 is seen to be stronger since Action-1 B_2-dominates Action-3. This is seen by the fact that although the accumulated graphs cross for different levels of accumulated belief, the area where Action-1 dominates is larger than for Action-3.

This way, the SIS entity is able to conclude that it should prefer Action-1 over the other two actions available. This decision is possible to do in an autonomous way even though there is a significant overlap in expected utilities and ordinary concepts of maximising the expected utility would not be able to discriminate between the various alternatives.

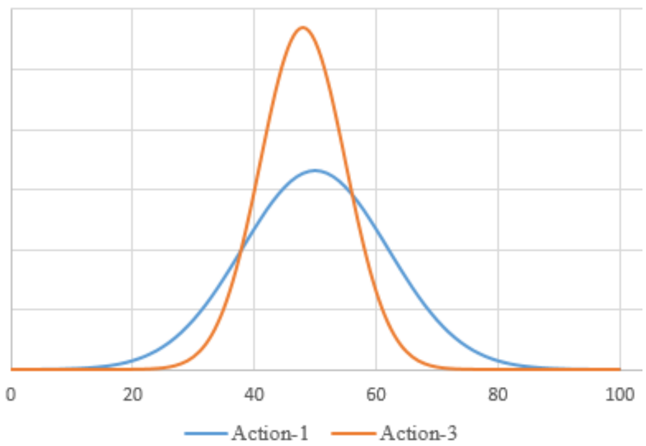

Figure 6. Comparing Action-1 and Action-3 (distribution of belief in expected values)

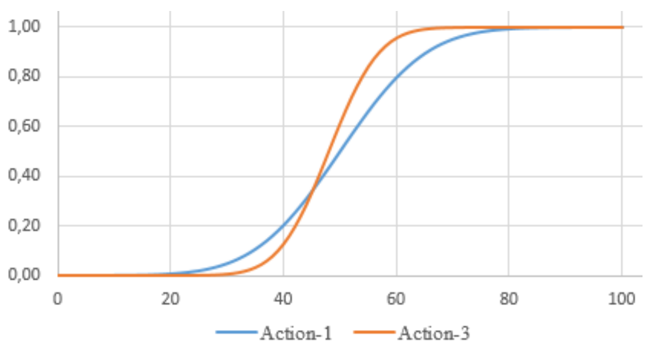

Figure 7. Comparing Action-1 and Action-3 (accumulated distribution of belief in expected values)[2]

7. CONCLUSIONS

In autonomous agent decision problems, it is usually impossible to assign precise numerical values to the different components of a decision model, and thus there is a need for representation and evaluation mechanisms that can effectively handle uncertain information in an autonomous way. Such models are not always immediately available, and to alleviate this problem, higher-order models such as belief mass representation can add both information and transparency, thus enabling a more discriminative analysis than, for example, using intervals alone are being proposed. In this paper, we describe a higher-order framework based on an evaluation method using a belief mass interpretation of the data involved. It involves a model where second-order information (in the sense of belief distributions over probabilities and values as well as aggregations thereof) is used for analysing decision trees. The framework presented in this paper is put to use for use in SIS decision situations where autonomy dictates the possible decision reasoning.

8. ACKNOWLEDGEMENTS

This research was supported by the EU-project Co-Inform (Co-Creating Misinformation-Resilient Societies H2020-SC6-CO-CREATION-2017) and strategic grants (SFO) from the Swedish government within ICT – The Next Generation.

9. REFERENCES

[1] D. Dubois, Representation, Propagation, and Decision Issues in Risk Analysis Under Incomplete Probabilistic Information, Risk Analysis **30**(3) (2010) 361–368.

[2] M. Danielson and L. Ekenberg, Computing Upper and Lower Bounds in Interval Decision Trees, European Journal of Operational Research **181** (2007) 808–816.

[3] J. Rohmer and C. Baudrit, The use of the possibility theory to investigate the epistemic uncertainties within scenario-based earthquake risk assessments, Natural Hazards **56**(3) (2010) 613–632.

[4] A.F. Shapiro and M.C. Koissi, Risk Assessment Applications of Fuzzy Logic. Casualty Actuarial Society, Canadian Institute of Actuaries, Society of Actuaries (2015).

[5] P. Dutta, Human Health Risk Assessment Under Uncertain Environment and Its SWOT Analysis, The Open Public Health Journal **11** (2018) 72–92.

[6] M. Danielson, L. Ekenberg and A. Larsson, DecideIT 3.0: Software for Second-Order Based Decision Evaluations, Bock, De J; de Campos, C; de Cooman, G; Quaeghebeur, E; Wheeler, G (Ed.): ISIPTA'19: Proceedings of the Eleventh International Symposium on Imprecise Probability: Theories and Applications, Proceedings of Machine Learning Research **103** (2019) 121–124.

[7] M. Danielson, L. Ekenberg, and A. Larsson, Belief Distribution in Decision Trees, International Journal of Approximate Reasoning, **46**(2) (2007) 387–407.

[8] M. Danielson, L. Ekenberg, A. Larsson, and D. Sundgren, Second-Order Risk Constraints in Decision Analysis, Axioms **3** (2014) 31–45.

[9] S. Kotz and J.R. van Dorp, Beyond Beta, Other Continuous Families of Distributions with Bounded Support and Applications, World Scientific Press, Singapore (2004).

[10] L. Ekenberg and J. Thorbiörnson, Second-Order Decision Analysis, International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, **9**(1) (2001) 13–38.

[11] L. Ekenberg, M. Danielson, and J. Thorbiörnson, Multiplicative Properties in Decision Trees, International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems **14**(3) (2006) 293–316.

[12] D. Sundgren, M. Danielson, and L. Ekenberg, Warp Effects on Calculating Interval Probabilities, International Journal of Approximate Reasoning, **50**(9) (2009) 1360–1368.

[2] Note that the figures are simplified in the sense that the belief distributions displayed in the example are symmetric. In real-life situations, the belief in the resulting expected values are almost invariably asymmetric due to calculation aspects as discussed in [12].