COMMON-SENSE SUPPORT FOR AUTONOMOUS AND ROBUST DECISION MAKING

Lionel Daniel and Valérie Roy

Applied Mathematics Centre, Mines ParisTech, 06904 Sophia Antipolis Cedex, France

ABSTRACT

Future space science mission will involve unmanned spacecrafts performing in hazardous environment at far distance from Earth. We therefore theoretically address the problem of making autonomous and robust decisions wrt inconsistent and uncertain information. To achieve such decisions, we propose to equip a spacecraft with the paraconsistent probabilistic reasoning, which is a new technique to infer information from inconsistent and probabilistic knowledge bases. This technique satisfies several principles indented to define the *common sense*. We also propose to design the programmed spacecraft behaviours in a synchronous language; such language are utilised to develop, verify, and certify safety-critical embedded system. By injecting some common sense into decision systems, we hope to make them more trustworthy.

Key words: common sense, decision making, paraconsistent probabilistic reasoning.

1. INTRODUCTION

The success of future space missions will rely on the spacecraft aptitude for making reliable decisions. In this paper, we thus propose a methodology for onboard decision making, focused on spacecraft autonomy. This theoretical methodology, depicted in Fig. 1, is twofold. On earth, space engineers specify the deterministic spacecraft behaviours. Aboard, these uploaded behaviours conduct activities according to sensory data and some *common sense*.

After sketching the spacecraft behaviours programming performed by engineers on Earth, we introduce the spacecraft decision process that manages these behaviours aboard. Then, in section 2, we provide a formalisation for the sensory data: the knowledge base. Finally, we define the paraconsistent probabilistic reasoning as a set of tools for knowledge bases, and we argue that they provide autonomous and robust decision making.

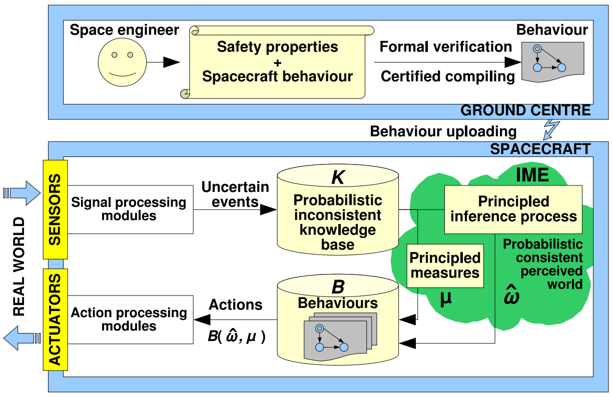

Figure 1. Methodology for robust decision making.

1.1. Spacecraft behaviours design

Firstly, on earth, space engineers specify the deterministic spacecraft behaviours. Behaviours represent tasks to realise wrt the current situation. In the following example, behaviour b_3 executes sub-behaviours b_1 and b_2 conditionally to c_1, which depends on the probability that event e_1 occurs in the current situation:

e_1: "*camera-1 detects life on Mars*"
b_1: "*inform ground centre about the probability of* e_1"
b_2: "*focus camera-2 on camera-1's target*"
c_1: "*probability of* e_1 *is higher than 80%*"
b_3: "*if* (c_1) *then (suspend low priority behaviours and execute simultaneously* b_1 *and* b_2*; when* b_1 *and* b_2 *terminates, resume low priority behaviours)*"

Notice that behaviour b_3 is deterministic iff c_1 is either **t**rue or **f**alse, ie iff the probability of e_1 is computable.

A behaviour, together with its set of safety properties, is written in a synchronous programming language. Such "languages have been designed to allow the unambiguous description of reactive, embedded real-time systems. The common foundation for these languages is the synchrony hypothesis, which considers computations to not take any time. This abstraction allows to separate the concerns functionality and real-time characteristics, and thus facilitates the design of complex embedded systems"[1]. In this paper, we propose to use SCADE Suite[2], which is an Inte-

[1] This description is excerpted from the SYNCHRON'2009 workshop website: `http://www.dagstuhl.de/09481`

[2] SCADE Suite is a trademark of Esterel Technologies SA. All rights

Proc. 'IJCAI–09 Workshop on Artificial Intelligence in Space', Pasadena, California, US
17–18 July 2009 (ESA SP-673, September 2009)

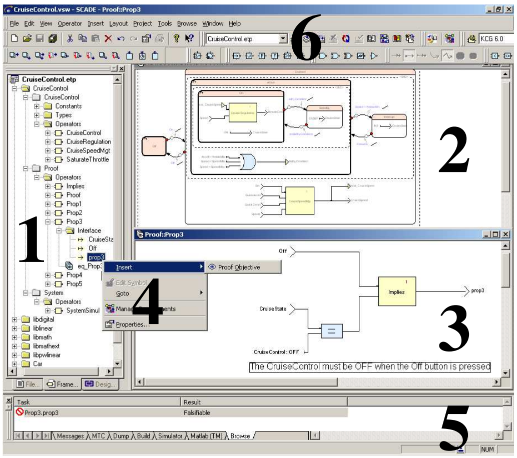

The CruiseControl must be OFF when the Off button is pressed

Task	Result
Prop3.prop3	Falsifiable

Figure 2. *SCADE Suite screenshot showing the design, the verification, and the compilation of a behaviour.*

grated Development Environment to design, verify, then generate certified[3] code. It provides graphical and textual formal languages, both with data-flow and control-flow synchronous programming styles. These languages comprise instructions to modularly parallelise, sequentialise, suspend, resume, and abort behaviours. The SCADE Suite screenshot in Fig. 2 shows **1**) the list of nodes, where a node represents a behaviour or a property, **2**) a behaviour, written in both data-flow (yellow blocks) and control-flow (blue and pink blocks), **3**) a property that a behaviour should satisfy, **4**) the model checking of the property, **5**) the result of the verification, and **6**) the behaviour compilation. The forbidden sign at the bottom left corner indicates that the behaviour does not satisfy the property. In which case, a scenario leading to the violation of the property is generated, helping thus engineers to debug the behaviour. Finally, the behaviour is uploaded aboard the spacecraft into a repository called B.

1.2. Behaviours driven by common sense

Once onboard, behaviours B determine the spacecraft decisions, wrt the current situation depicted by sensors. Because of the hazardous spacecraft environment, sensory data are tainted with uncertainty; eg, the processing of the *camera-1* images could lead to uncertain event, eg “*probability of e_1 is lower than 30%*”; such events may be imprecise due to missing or partial sensory data resulting from sensor failure or power loss. Uncertain events are stored into a knowledge base K, which tends to be inconsistent due to the multi-sensor context. Thus, the spacecraft must act wrt an imprecise, uncertain, and inconsistent knowledge base. In section 3.1, we will propose a process, called IME, that infers from K a precise and probabilistically consistent world model $\hat{\omega}$. In addition to IME, we will define several principled measures μ for knowledge bases. These measures enable engineers to specify behaviours such as “*if* $(\mu^{\text{impr}}(K_{\text{camera-1}}) \geq 80\%)$ *then (execute* b_2*)*”, which commands to the camera-2 to focus on camera-1’s target when camera-1 provides the spacecraft with too imprecise data. Thus, the spacecraft actions are computed by evaluating behaviours B wrt to IME(K) and the measures: $B(\hat{\omega}, \mu)$. The key for autonomous and robust decision making resides in the *common sense* underlying IME and μ.

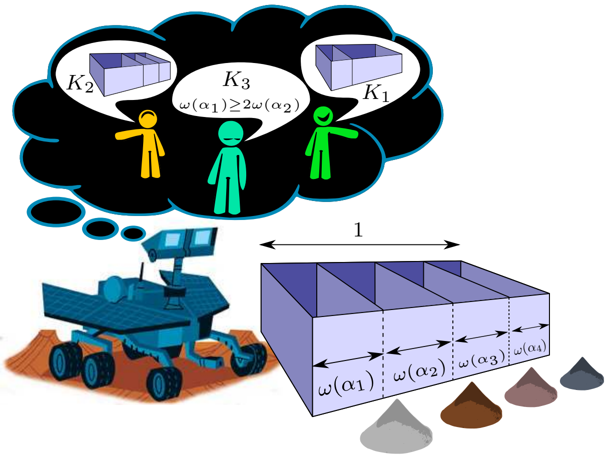

Figure 3. *Storage box with sliding walls; the compartments capacity is adjusted to the amount of soils.*

2. PROBABILISTIC KNOWLEDGE REPRESENTATIONS

In this section, we define the probabilistic language, named $\mathbb{K}$, in which the uncertain events are expressed. These events constrain a probability distribution on sentences of a propositional language. Then we define a knowledge base as a set of such constraints. We also define the consistency of a knowledge base as the satisfiability of its set of constraints. Finally, we introduce a more general knowledge representation to deal with inconsistencies. But first, we motivate our choice towards probabilistic reasoning through voting theory.

2.1. Motivating example: voting theory

Voting theory is a theory of electing a societal preference from individual preferences. In the following example, a rover will have to achieve a consensus about resource allocation from the possibly conflicting preferences of its embedded agents.

Suppose a rover is scouting a surface for soil sampling. This rover embeds several scientific agents, ie computer programs, that together decide on the amount of each soils to carry back to the main station where further analysis will be performed. During the journey, the rover stows the soils in a storage box having sliding walls (see

reserved. See `http://www.esterel-technologies.com/`

[3] Code generation qualifiable for DO-178B up to Level A, certifiable for IEC 61508 certified at SIL 3 and EN 50128 certified at SIL 3/4.

Fig. 3): this allows to adjust the capacity of each compartment to the amount of a collected soil. The agents have different interests, eg, one agent focuses on organic chemistry, whereas another agent focuses on rare soils. Due to the finite capacity of the storage box, these interests may be conflicting, eg, the latter agent may want to carry back a maximum amount of a rare soil, even if this soil is much less inspiring from the organic standpoint than an abundant soil. Notwithstanding the possibly conflicting agents' interests, the rover must achieve a consensus about the capacity of each compartment of the storage box; we formally state this problem as follows.

A rover embedding $I \in \mathbb{N}$ agents stows soil samples in $J \in \mathbb{N}$ compartments $\{ \alpha_1, \alpha_2, \ldots, \alpha_J \}$ of a box with sliding walls. A space distribution ω is a function that maps each compartment to its capacity and satisfies these two assumptions: (A1) the box volume is 1 cubic decimetre, ie $1 = \sum_{j=1}^{J} \omega(\alpha_j)$, and (A2) each compartment capacity is positive, ie $\forall j \in \{ 1, 2, \ldots, J \}, \omega(\alpha_j) \geq 0$.

Each agent i expresses a set K_i of wishes for the space distribution ω. For example, i may wish to allocate at least twice more space to soil α_1 than to soil α_2, ie $\omega(\alpha_1) \geq 2 * \omega(\alpha_2)$, and may wish that the total amount allocated to α_1 and α_2 is within 0.2 and 0.3 cubic decimetre, ie $0.2 \leq \omega(\alpha_1) + \omega(\alpha_2) \leq 0.3$; definition 3 will formalise the wishes. Besides, the rover affixes to each agent i a reliability degree $\sigma_i \in \mathbb{R}^+$, which tends towards 0 when the rover tends to consider i as reliable; eg, i will be labelled as reliable if, in case the rover had fulfilled its wishes without considering other agents' wishes, its wishes would have enabled a high science return.

Thus, the rover must implement a voting system $\mathcal{I}$ yielding the space distribution $\hat{\omega}$ that *best* conciliates the wishes K_i of each agent i, according their reliability σ_i and some *common sense*; formally, $\hat{\omega} \overset{\text{def}}{=} \mathcal{I}_{[\sigma_1,\ldots,\sigma_I]}(\cup_{i=1}^{I} K_i)$, where $\mathcal{I}$ must satisfy several principles intended to define the common sense. By interpreting assumptions (A1) and (A2) as the Kolmogorov's axioms for probability, space distributions can be identified as *probability* distributions (these terms will be defined in section 2.2). We therefore take the *probabilistic* standpoint to define $\mathcal{I}$ as an inference process, of which a definition will be given in section 3.1; this motivates us to study the paraconsistent probabilistic reasoning.

2.2. Probabilistic language

Let Θ be a propositional language generated by $\Theta ::= (\Theta \wedge \Theta) \mid (\Theta \vee \Theta) \mid (\neg \Theta) \mid var$, where *var* is a finite set of propositional variables being either **t**rue or **f**alse, and where logical connectives $\wedge$, $\vee$, and $\neg$ have their respective classical semantic[4] *and*, *or*, and *not*. These variables represent the application domain of the spacecraft sensors, like event e_1 in the introductory example of section 1.1. In the sequel, unless explicitly stated, Θ is supposed fixed. Also, the propositions, usually noted θ, ϕ, or ψ, are supposed belonging to Θ. Let $\models \theta$ denote a tautology θ. Furthermore, let $\alpha_\Theta \overset{\text{def}}{=} \{ \alpha_j \mid j = 1, 2, \ldots, J \}$ denote the set of minterms[5] of Θ, where $J \overset{\text{def}}{=} 2^{|var|}$ with $|var|$ being the number of propositional variables. Also, let $\alpha_\theta \overset{\text{def}}{=} \{ \alpha_j \mid \models (\neg\alpha_j \vee \theta) \}$ denote the minterms of a proposition θ. Finally, each proposition θ is supposed to be in the canonical disjunctive normal form, ie $\theta = \bigvee_{\alpha \in \alpha_\theta} \alpha$.

Definition 1. *Kolmogorov's axioms for probability are:*
(P1) if $\models \theta$ then $\omega(\theta) = 1$;
(P2) if $\models \neg(\theta \wedge \phi)$ then $\omega(\theta \vee \phi) = \omega(\theta) + \omega(\phi)$,
where ω is a function from Θ to [0:1], $\theta, \phi \in \Theta$.

Definition 2. *A probability distribution ω is a function that satisfies Kolmogorov's axioms for probability. We denote by Ω_Θ, or by Ω when it is unambiguous, the set of probability distributions underlain by a given propositional language Θ.*

Notice that the minterms of Θ are mutually exclusive, ie $\models \neg(\alpha_i \wedge \alpha_j)$ for any two distinct minterms α_i and α_j. Since θ is a disjunction of minterms, $\omega(\theta)$ equals $\omega(\bigvee_{\alpha \in \alpha_\theta} \alpha)$ by definition, and equals $\sum_{\alpha \in \alpha_\theta} \omega(\alpha)$ by axiom (P2). Thus, a probability distribution ω can be seen as a function from α_Θ to [0:1], hence as a point $[\omega(\alpha_1); \omega(\alpha_2); \ldots; \omega(\alpha_J)]$ in an Euclidean space of dimension J such that its j^{th} coordinate $\omega_j \in [0{:}1]$ equals $\omega(\alpha_j)$. Furthermore, [6, pages 13–15] shows that a point $\omega \in \mathbb{R}^J$ in an Euclidean space of dimension J denotes a probability distribution iff $\omega \geq \vec{0}$ and[6] $1 = \sum_{j=1}^{J} \omega_j$. Thus, writing $\Omega \subset \mathbb{R}^J$ makes sense.

In this paper, we identify knowledge with a possibly unsatisfiable multiset of constraints on a probability distribution ω. Each constraint c, ie each item of knowledge, is an inequality with the following general form: $b \geq f(\omega)$, where $b \in \mathbb{R}$ and $f : \mathbb{D} \mapsto \mathbb{R}$ such that $\Omega \subseteq \mathbb{D} \subseteq \mathbb{R}^J$. If f is a linear function, ie if c has the form $b \geq [a_1, a_2, \ldots, a_J] * \omega$, where a_j are real numbers such that $1 = \sqrt{\sum_{j=1}^{J} {a_j}^2}$, then c is said to be a *linear* constraint.

Definition 3 (Linear knowledge base)**.** *A linear knowledge base is a multiset of linear constraints. We denotes by $\mathbb{K}$ the set of linear knowledge bases.*

Notice that a linear knowledge base $K \in \mathbb{K}$ is simply a matrix inequality $K.B \geq K.A * \omega$, where $K.B$ and $K.A$

[4] Classical semantic

θ	ϕ	$\theta \wedge \phi$	$\theta \vee \phi$	$\neg\theta$
false	**f**alse	**f**alse	**f**alse	**t**rue
false	**t**rue	**f**alse	**t**rue	**t**rue
true	**f**alse	**f**alse	**t**rue	**f**alse
true	**t**rue	**t**rue	**t**rue	**f**alse

[5] A *minterm* is a sentence of a propositional language. A minterm has the form $\bigwedge_{v \in var} \pm v$, where *var* is the set of propositional variables and where $\pm v$ means either $\neg v$ or v.

[6] Remember the assumption (A1) $1 = \sum_{j=1}^{J} \omega(\alpha_j)$, and (A2) $\forall j \in \{ 1, 2, \ldots, J \}, \omega(\alpha_j) \geq 0$ in the motivating example about voting theory at section 2.1 on the previous page.

are defined as follows, and where $I \in \mathbb{N}$ is the number of constraints:

$$K.B \overset{\text{def}}{=} \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_I \end{bmatrix} \quad K.A \overset{\text{def}}{=} \begin{bmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,J} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,J} \\ \vdots & \vdots & \ddots & \vdots \\ a_{I,J} & a_{I,2} & \cdots & a_{I,J} \end{bmatrix}$$

Despite their simplicity, linear knowledge bases generalise widely used bases such as sets of propositions or sets of conditional probabilities.

Definition 4 (Models of K)**.** *A model of a knowledge base K is a probability distribution satisfying all the constraints in K. We denote by Ω_K and Ω_c the set of models of K and $\{c\}$, respectively, where c is a constraint.*

Definition 5 (Consistency)**.** *A knowledge base K is consistent iff $\Omega_K \neq \emptyset$, ie, iff there exists a probability distribution satisfying all the constraints in K; otherwise, K is said to be inconsistent.*

2.3. Blur probabilistic representation

In this section, we present a more general probabilistic representation of knowledge: a candidacy function. As complex numbers were a new representation of numbers to deal with negative square roots, we propose candidacy functions as a new representation of knowledge to deal with *paraconsistent* reasoning.

Definition 6. *A candidacy function C is a function from Ω to* $[0{:}1]$ *such that $C(\omega) = 1$ means ω is a candidate for representing the real world.*

Definition 7 (Best candidates wrt C)**.** *The non-empty set of probability distributions that are the best candidates for representing the real world, wrt a candidacy function C, is defined as follows:*

$$\hat{\Omega}_C \overset{\text{def}}{=} \arg\max_{\omega \in \Omega} C(\omega)$$

In this paper, we only consider candidacy functions satisfying the following principle called *concession*, which talks about knowledge fusion, ie merging two candidacy functions.

ⓐ *Concession*: A candidacy function C is conceding iff, for any probability distribution ω not dismissed by C' from representing the real world, the candidacy function $C * C'$, resulting from the merging of C with another candidacy function C', does not dismiss ω.

$$\text{if } C'(\omega) > 0 \text{ then } (C * C')(\omega) > 0$$

We do not give further insights into candidacy functions in this paper because our focus is on linear knowledge bases. We nevertheless provide the construction of a candidacy function C_K corresponding to a given knowledge base K. We require of C_K to satisfy several properties. Firstly, $C_K(\omega) = 1$ iff ω is a model of K, otherwise, $C_K(\omega)$ is intended to represent the degree of satisfiability of K. Secondly, since K is a multiset of constraints, we assume $C_K(\omega) = \prod_{c \in K} C_{\{c\}}(\omega)$, ie C_K only depends on the degree of satisfiability of each constraint $c \in K$. Thirdly, remember that a linear constraint c has the form $b \geq [a_1, a_2, \ldots, a_J] * \omega$. Also, notice that equality $b = [a_1, a_2, \ldots, a_J] * \omega$ denotes a hyperplane in an Euclidean space of dimension J. This hyperplane separates the probability distributions satisfying c from those that do not. We thus suppose that, from two probability distributions not satisfying c, the one most satisfying c is the closest to the hyperplane, wrt the Euclidean distance. We define the normalised signed Euclidean distance between a probability distribution ω and the hyperplane of c as follows: $\mathcal{D}_c^\omega \overset{\text{def}}{=} (b - [a_1, a_2, \ldots, a_J] * \omega) * \sqrt{J}$. Finally, we interpret c as a random polynomial with a random variable b whose cumulative distribution function h_{σ_c} is half-Gaussian with a standard deviation $\sigma_c \in \mathbb{R}^+$:

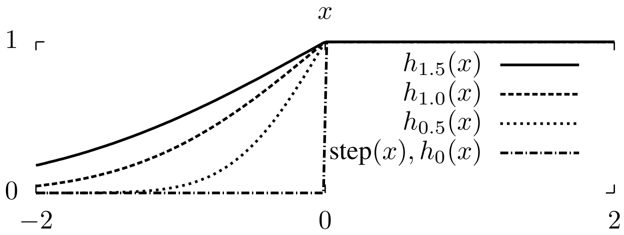

Figure 4. Half-Gaussian cumulative distribution function

$$C_K(\omega) \overset{\text{def}}{=} \prod_{c \in K} h_{\sigma_c}(\mathcal{D}_c^\omega)$$

where h_σ is defined as follows:

$$h_\sigma(x) \overset{\text{def}}{=} \begin{cases} 1 + \text{erf}\left(\frac{x}{\sigma\sqrt{2}}\right) & \text{if } \sigma > 0 \text{ and } x < 0, \\ \text{step}(x) & \text{otherwise.} \end{cases}$$

and where the error (erf) and the step (step) functions are respectively defined as follows:

$$\text{erf}(x) \overset{\text{def}}{=} \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt \qquad \text{step}(x) \overset{\text{def}}{=} \begin{cases} 1 & \text{if } x \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Notice that graphs of functions step and h_σ are drawn in Fig. 4, for several values of standard deviation σ.

We now stress two facts. Firstly, $\hat{\Omega}_{C_K} = \Omega_K$ when K is consistent. Secondly, the construction of C_K is *language invariant*, ie, introducing a new variable v into the underlying propositional language Θ of a knowledge base $K \in \mathbb{K}$, yields a new language Θ' underlying a new knowledge base $K' \in \mathbb{K}$ such that $\forall \omega' \in \Omega_{K'}, C_{K'}(\omega') = C_K(\omega)$, where $\omega \in \Omega_K$ is such that $\forall \alpha \in \alpha_\Theta, \omega(\alpha) \overset{\text{def}}{=} \omega'(\alpha)$; this is due to coefficient $\sqrt{J}$ in $\mathcal{D}_c^\omega$.

3. PARACONSISTENT PROBABILISTIC REASONING

In this section, we define five tools for the paraconsistent probabilistic reasoning. We then propose to utilise them for achieving autonomous and robust decision making.

3.1. Inference process and measures

Let $K, K' \in \mathbb{K}$ be two linear knowledge bases; we suppose they are underlain by a common propositional language, but this is not a restriction since the next tools all satisfies *language invariance*. Let $\hat{\omega}_K \in \hat{\Omega}_{C_K}$ be a best candidate for representing the real world, wrt C_K. Our five tools for the paraconsistent probabilistic reasoning are:

- an *inference* process:
  $$\text{IME}(K) \overset{\text{def}}{=} \arg \max_{\hat{\omega}_K \in \hat{\Omega}_{C_K}} E(\hat{\omega}_K)$$
  where E is the entropy of a probability distribution:
  $$E(\omega) \overset{\text{def}}{=} -\sum_{j=1}^{J} \omega_j * \log(\omega_j)$$
- a *dissimilarity* measure:
  $$\mu^{\text{dis}}_{\mathcal{L}_\infty}(K, K') \overset{\text{def}}{=} \max_{\omega \in \Omega} |C_K(\omega) - C_{K'}(\omega)|$$
- an *inconsistency* measure:
  $$\mu^{\text{icst}}(K) \overset{\text{def}}{=} 1 - C_K(\hat{\omega}_K)$$
- together with its *culpability* measure:
  $$\mu^{\text{icst}}(c \in K) \overset{\text{def}}{=} 1 - C_c(\hat{\omega}_K)$$
  which quantifies the amount of inconsistency brought to a knowledge base K by an item c belonging to K;
- an *imprecision* measure:
  $$\mu^{\text{impr}}(K) \overset{\text{def}}{=} \mathcal{V}(\hat{\Omega}_{C_K})$$
  where $\mathcal{V}(\hat{\Omega}_{C_K})$ is the *volume* filled by the set $\hat{\Omega}_{C_K}$; informally, $\mathcal{V}$ intends to compute the number of probability distributions that best represent the real world wrt C_K: the greater $\mathcal{V}(\hat{\Omega}_{C_K})$, the more imprecise C_K;
- an *incoherence* measure:
  $$\mu^{\text{icoh}}(K, K') \overset{\text{def}}{=} C_K(\hat{\omega}_K) * C_{K'}(\hat{\omega}_{K'}) - C_{K \cup K'}(\hat{\omega}_{K \cup K'})$$
  which is roughly the distance between the maximal consistency $C_K(\hat{\omega}_K) * C_{K'}(\hat{\omega}_{K'})$ and the real consistency $C_{K \cup K'}(\hat{\omega}_{K \cup K'})$ between K and K';

These five tools are principled, ie they satisfy principles. We now present a selection of such principles that all together are intended to define the *common sense*.

3.2. Common sense

3.2.1. Notions of convergence

The notion of converging sequence of knowledge bases is founded upon a metric: the Hausdorff distance.

Definition 8 (Hausdorff distance)**.** *The Hausdorff distance $\mathcal{H}$ of two non-empty compact (bounded and closed) sets X and Y of points in a Euclidean space is defined as follows, where $\mathcal{L}_1(x, y)$ denotes the Euclidean distance between points x and y:*

$$\mathcal{H}(X, Y) \overset{\text{def}}{=} \inf \left\{ \delta \;\middle|\; \text{and} \begin{array}{l} \forall x \in X, \exists y \in Y, \delta \geq \mathcal{L}_1(x, y) \\ \forall y \in Y, \exists x \in X, \delta \geq \mathcal{L}_1(x, y) \end{array} \right\}$$

Jeff Paris argues in [6, pages 89–91] that $\mathcal{H}(\Omega_{K_1}, \Omega_{K_2})$ corresponds to the distance between knowledge contents when K_1 and K_2 are consistent knowledge bases (and underlain by the same propositional language). Thus, a sequence of consistent knowledge bases $\kappa\colon \mathbb{N} \mapsto \mathbb{K}$ converges to a consistent knowledge base $K \in \mathbb{K}$, denoted by $\lim^{\mathcal{H}}_{i \to \infty} \kappa_i = K$, iff $\forall \varepsilon \in \mathbb{R}^+, \exists N_\varepsilon \in \mathbb{N}, \forall i \geq N_\varepsilon, \mathcal{H}(\Omega_{\kappa_i}, \Omega_K) < \varepsilon$. However, when dealing with inconsistent knowledge bases, Ω_{K_1} or Ω_{K_2} are empty, hence, $\mathcal{H}$ can not compute their distance. Therefore, we use $\hat{\Omega}_{C_K}$ instead of Ω_K and we define a new metric $\mu^{\text{dis}}_{\mathcal{H}}(K_1, K_2)$ that is equivalent to $\mathcal{H}$ when dealing with consistent knowledge bases, where $[\omega \sqcup r]$ denotes the vertical concatenation of a vector ω with a scalar $r \in \mathbb{R}$:

$$\mu^{\text{dis}}_{\mathcal{H}}(K_1, K_2) \overset{\text{def}}{=} \mathcal{H}\left(\begin{array}{l} \left\{ [\omega \sqcup C_{K_1}(\omega)] \;\middle|\; \omega \in \hat{\Omega}_{C_{K_1}} \right\}, \\ \left\{ [\omega \sqcup C_{K_2}(\omega)] \;\middle|\; \omega \in \hat{\Omega}_{C_{K_2}} \right\} \end{array} \right)$$

We consider that the agent's knowledge has two levels: an internal level (the agent's epistemic state represented by a candidacy function C) where the knowledge management is performed (eg: to resolve the inconsistency), and an external level (the agent's visible knowledge, or exposed state, represented by $\hat{\Omega}_C$) where a higher kind of reasoning is realised (eg: to use an inference process). We employ $\mu^{\text{dis}}_{\mathcal{L}_\infty}$ to measure the distance between two internal levels of knowledge, whereas we use $\mu^{\text{dis}}_{\mathcal{H}}$ to compare two external levels of knowledge. We thus obtain two notions of convergence, one for each level.

Definition 9 (Convergence wrt μ^{dis})**.** *A sequence $\kappa\colon \mathbb{N} \mapsto \mathbb{K}$ converges to a knowledge base $K \in \mathbb{K}$ wrt a metric μ^{dis}, denoted by $\lim^{\mu^{\text{dis}}}_{i \to \infty} \kappa_i = K$, iff $\forall \varepsilon \in \mathbb{R}^+, \exists N_\varepsilon \in \mathbb{N}, \forall i \geq N_\varepsilon, \mu^{\text{dis}}(C_{\kappa_i}, C_K) < \varepsilon$.*

3.2.2. Common-sensical principles

In this paper, we adhere to the proposal that reasoning common-sensically is reasoning while applying intuitive principles; most of these principles are underlain by ideas coming from [6, Chapter 7]. We now present several principles satisfied by our inference process and measures.

Principles for an inference process $\mathcal{I}$ In order to legibly define principles ⓓ to ⓙ, we shall assume that $\mathcal{I}$ satisfies *uniqueness* and is *language invariant*; this latter property is a consequence of *irrelevant information* principle (see ⓒ then take K_2.beliefs $= \emptyset$). Thus, when declaring some knowledge bases underlain by some propositional variables, eg $K_1 \in \mathbb{K}_{\text{vars}_1}$ and $K_2 \in \mathbb{K}_{\text{vars}_2}$, we can freely substitute their respective underlying language by a common one that possesses all their propositional variables, eg $K_1, K_2 \in \mathbb{K}_{\text{vars}_1 \cup \text{vars}_2}$, therefore, we can omit the reference to their language, eg $K_1, K_2 \in \mathbb{K}$.

ⓑ *Uniqueness & determinism*: An inference process should deterministically return a unique probability distribution.

ⓒ *Irrelevant information* (Extension of [6, page 87]): Entirely irrelevant information should be ignored by an inference process. Let $K_1 \in \mathbb{K}_{\text{vars}_1}$ and $K_2 \in \mathbb{K}_{\text{vars}_2}$ be two knowledge bases underlain by two disjoint sets of propositional variables, ie vars$_1 \cap$ vars$_2 = \emptyset$. Let θ_1 be a propositional sentence underlain by vars$_1$. Then K_2 is said to be irrelevant to K_1 and θ_1.

$$(\mathcal{I}(K_1))(\theta_1) = (\mathcal{I}(K_1 \cup K_2))(\theta_1)$$

ⓓ *Equivalence* (Extension of [6, page 82]): Equal information should be inferred from equivalent knowledge bases.

$$\text{if } \mu_{\mathcal{H}}^{\text{dis}}(K_1, K_2) = 0 \text{ then } \mathcal{I}(K_1) = \mathcal{I}(K_2)$$

ⓔ *Renaming* (Due to [6, page 95]): An inference process should be insensitive to a renaming of the propositional variables. For any knowledge base $K \in \mathbb{K}$ underlain by n propositional variables, let π be a permutation over the natural numbers in $[1{:}2^n]$. Furthermore, when applied to K, π is defined by $\pi(K).A_{r,c} \overset{\text{def}}{=} K.A_{r,\pi(c)}$ and $\pi(K).b \overset{\text{def}}{=} K.b$ where r and c denote the r^{th} row and the c^{th} column of the matrix $K.A$. When applied to a vector ω, π is defined by $\pi(\omega)_j \overset{\text{def}}{=} \omega_{\pi(j)}$ where j denotes the j^{th} element of vector ω.

$$\mathcal{I}(\pi(K)) = \pi(\mathcal{I}(K))$$

ⓕ *Obstinacy* (Extension of [6, page 90]): Additional support for what is already known should be ignored by an inference process.

$$\text{if } \mathcal{I}(K_1) \in \hat{\Omega}_{K_2} \text{ then } \mathcal{I}(K_1) = \mathcal{I}(K_1 \cup K_2)$$

ⓖ *Independence* (Due to [6, page 101]): The absence of any information linking two events should be identified with the conditional independence; justifications for this principle are given in [8]. Let L be a propositional language having three variables v_1, v_2, and v_3, and underlying a knowledge base K defined as follows, with $a, b \in [0{:}1]$:

$$K.\text{beliefs} \overset{\text{def}}{=} \left\{ \begin{array}{r} \omega(v_1) = a \\ \omega(v_2 \mid v_1) = b \\ \omega(v_3 \mid v_1) = c \end{array} \right\}$$

The *independence* principle states that v_2 and v_3 should be treated as conditionally independent given v_1:

$$(\mathcal{I}(K))(v_2 \wedge v_3 \mid v_1) = b * c$$

ⓗ *Continuity* (Extension of [6, page 89]): Microscopic changes in the knowledge base should not cause macroscopic changes in the inferred information. This principle ensures a certain robustness in face of minor fluctuations in the knowledge base.

$$\text{if } \lim_{i\to\infty}^{\mu_{\mathcal{H}}^{\text{dis}}} \kappa_i = K \text{ then } \lim_{i\to\infty} \mathcal{I}(\kappa_i) = \mathcal{I}(K)$$

ⓘ *Open-mindedness* (Extension of [6, page 95]): An inference process should give the benefit of the doubt; this principle is a kind of precautionary principle. Let θ be any sentence of the underlying propositional language of a knowledge base K.

$$\text{if } \exists \omega \in \hat{\Omega}_K, \omega(\theta) > 0 \text{ then } (\mathcal{I}(K))(\theta) > 0$$

ⓙ *Relativisation* (Due to [6, page 100]): The probabilities an inference process would give if some event occurred should only depend on the knowledge conditioned by the occurrence of this event. For the sake of elegance of the following definitions, we express the knowledge bases in a non-normalised form. Let $K, K_1, K_2 \in \mathbb{K}$, $a_{ij}, b_i, a'_{ij}, b'_i, c \in \mathbb{R}$, $k, k', l_i, l'_i \in \mathbb{N}$, $\theta, \theta_i, \theta'_i, \varphi \in \Theta$, and $K \overset{\text{def}}{=} \{\, c = \omega(\varphi) \,\}$ with $0 < c < 1$.

$$K_1 \overset{\text{def}}{=} \left\{ b_i = \sum_{j=1}^{l_i} a_{ij} * \omega(\theta_i \mid \varphi) \;\middle|\; i = 1, \ldots, k \right\}$$

$$K_2 \overset{\text{def}}{=} \left\{ b'_i = \sum_{j=1}^{l'_i} a'_{ij} * \omega(\theta'_i \mid \neg\varphi) \;\middle|\; i = 1, \ldots, k' \right\}$$

Notice that K_1 expresses knowledge relatively to the occurrence of φ, whereas K_2 expresses knowledge in case φ does not occur. Thus, the *relativisation* principle states that the probability of θ given φ should only depend on $K \cup K_1$, when $K \cup K_1 \cup K_2$ is consistent:

$$\begin{array}{l} \text{if } \Omega_{K \cup K_1 \cup K_2} \neq \emptyset \\ \text{then } \mathcal{I}(K \cup K_1)(\theta \mid \varphi) = \mathcal{I}(K \cup K_1 \cup K_2)(\theta \mid \varphi) \end{array}$$

The demonstration of the following characterisation theorem appears in [7], then has been generalised to consistent polynomial knowledge bases in [9].

Theorem 1 ([6, theorem 7.9]). *When dealing with a* consistent *linear knowledge base,* ME $\overset{\text{def}}{=} \arg\max_{\omega \in \Omega_K} E(\omega)$ *is the unique inference process satisfying principles* ⓑ *to* ⓙ.

Theorem 2. IME(K) = ME(K) *when* $K \in \mathbb{K}$ *is consistent, and* IME *satisfies principles* ⓑ *to* ⓙ.

Principles for measure We only present two key principles for measures: language invariance and continuity.

ⓚ *Language invariance*: A measure μ of a knowledge base K is language invariant iff adding a variable into the underlying propositional language of K does not change its value. Similarly, if μ takes two arguments K_1 and K_2, then adding a variable into the underlying propositional language of K_1 and K_2 does not change its value. More formally, let $K_1, K_2 \in \mathbb{K}_{\text{vars}}$ be two knowledge bases underlain by a non-empty set of propositional variables vars, and let $K_1', K_2' \in \mathbb{K}_{\text{vars}\cup\{v\}}$ be the same bases as K_1 and K_2 with one variable $v \notin$ vars added into their underlying propositional language.

$$\mu(K_1) = \mu(K_1') \qquad \text{or} \qquad \mu(K_1, K_2) = \mu(K_1', K_2')$$

ⓛ *Continuity*: Microscopic changes in the knowledge base should not cause macroscopic changes in the measure. This principle ensures a certain robustness in face of minor fluctuations in the knowledge base.

$$\begin{array}{ll} \text{if} & \lim_{i\to\infty}^{\mu^{\text{dis}}_{\mathcal{L}^\infty}} \kappa_i = K \\ \text{then} & \lim_{i\to\infty} \mu(\kappa_i) = \mu(K) \\ \text{or} & \lim_{i\to\infty} \mu(\kappa_i, K') = \mu(K, K') \end{array}$$

3.3. Autonomous and robust decision making

IME satisfies principles ⓑ to ⓙ ensuring:

- autonomy, ie decisions are taken without recourse to humans: *uniqueness* (see ⓑ) ensures that an evaluation of a condition in a behaviour (see condition c_1 in the example at section 1.1 on page 1) always returns either **t**rue or **f**alse;
- determinism, ie decisions are explainable: *determinism* (see ⓑ) ensures that IME does not use any random function, hence the evaluation of the conditions in the behaviours are deterministic, therefore, the whole behaviour is deterministic;
- robustness, ie decisions are robust against slight fluctuations of sensory data: *continuity* (see ⓗ) ensures the continuity of IME, but a value of a condition in a behaviour can wobble. If this effect is undesirable, an engineer could design more sophisticated behaviours which compute spacecraft actions by applying some continuous functions to IME(K) (so that the spacecraft action continuously depends on the sensory data), but in which case, the model checker may not be able to formally verify the behaviour;
- fairness, ie IME equally trusts, or fairly conciliates, each uncertain event in K: this property is ensured by the knowledge formalisation, because K is a multiset, and by *concession* (see ⓐ), which avoid knowledge pieces to be ignored, even when they are inconsistent;
- backwards compatibility, ie the spacecraft decisions are not influenced by the addition of new sensors if these sensors provide data on new topics (hence an decision taken before the spacecraft upgrade is still valid): this is ensured by *irrelevant information* (see ⓒ).
- semantic analysis, ie decisions depend on the meaning of K, not on the syntax: the knowledge normalisation and the other principles are intended to make an inference process syntax invariant (eg, see ⓓ and ⓔ).

In addition to IME, we propose in section 3.1 several principled measurements of K. These measurements allow engineers to define behaviours that establish strategies for:

- mission planning, by measuring the *incoherence* between the current situation and the mission target, both described in terms of knowledge bases. A behaviour could be "if the mission target is too incoherent from the current situation depicted by the sensors, then the spacecraft should select a more achievable target".
- tackling unexpected events, by measuring the *dissimilarity* between the current situation and an expected one, both described in terms of knowledge bases;
- self-healing, by measuring the *culpability* of each sensor for making K inconsistent: spacecraft may decide to check then repair such a sensor.
- sensors recalibration, by measuring the *imprecision* and the *redundancy* of sensory data. For example, an exploring spacecraft may decide to widen its sensor coverage by decreasing the overlap of each sensor coverage, ie by increasing the dissimilarity between sensory data. However, when the spacecraft detects an interesting event, it may decide to focus its sensors on this event by increasing the overlap of each sensor coverage.

3.4. Computational complexities

In the sequel, we denote by m the number of inequalities in a knowledge base $K \in \Theta$, and by n the number of propositional variables. The space complexity of our knowledge representation is exponential wrt n. Besides, the time complexity of our inference process IME depends on the space complexity. Thus, in order to make IME tractable, we are investigating techniques that exponentially reduce the space complexity, like those in [2].

Space and time complexity The naive space complexity of our knowledge representation is $\mathcal{O}(m * 2^n)$. The following partitioning technique exponentially reduces this complexity. A knowledge base can be partitioned into p sub-bases of inequalities such that each sub-base does not share any propositional variable with the others. Notice that the knowledge in a partition is irrelevant to the knowledge in another partition. This partitioning technique is legitimate for any inference process satisfying principle ⓒ, like IME. Hence, the space complexity of the partitioned knowledge is only $\sum_{i=1}^{p} \mathcal{O}(m_i * 2^{n_i})$, with $n = \sum_{i=1}^{p} n_i$ and $i = 1, \ldots, p$ where m_i and n_i are respectively the number of inequalities and propositional variables of the i^{th} partition. Due to the partitioning, IME applied to a knowledge base computes $p * 2$ optimisations over 2^{n_i} variables within $[0{:}1]$ instead of two optimisations over 2^n variables. If p is large then $2^{n_i} \ll 2^n$, and $\sum_{i=1}^{p} \mathcal{O}(m_i * 2^{n_i})$ might become a tractable space complexity.

The time complexity of IME(K) relies on the time complexity for maximising p times the function E over $\hat{\Omega}_{C_{K_i}}$, which is an maximisation of C_{K_i} over 2^{n_i} variables with $i = 1, \ldots, p$, where K_i is a partition of K. We know that C_{K_i} is not only continuous and log-concave but also non-smooth (see the non-smoothness of h_σ in Fig. 4). Thus, the time complexity of $\hat{\Omega}_{C_{K_i}}$ is the same as maximising a concave non-smooth function over the convex set $[0{:}1]^{2^{n_i}}$ constrained by the linear equality $1 = \sum_{j=1}^{2^{n_i}} \omega_j$ (see [10]).

On bounding and approximating techniques In addition, there exist techniques to smooth out a log-concave function (see [4]) enabling us to not only use faster optimisation algorithms (see [3]), but to also compute a hat function (see [1]) that allows arbitrarily precise approximation of $\hat{\Omega}_{C_{K_i}}$. Furthermore, an easier-to-compute entropy function is proposed in [5], which accelerates each function evaluation during the optimisation process.

Tractable consensus decision making In section 2.1, we propose to use IME for computing a consensus among the agents about the capacity of the J compartments, where $J \in \mathbb{N}$ must be a power of two. In this situation, the space complexity of a knowledge base K containing m agents' wishes is $\mathcal{O}(m * J)$. Thus, IME(K) may be tractable.

4. CONCLUSION

In this paper, we introduce the paraconsistent probabilistic reasoning as the only solution to a certain kind of consensus decision making (see section 2.1), and a possible solution to autonomous and robust decision making (see section 3.3). We introduce the paraconsistent probabilistic reasoning as a set of principled tools (measures and inference process, see 3.1) that deal with possibly inconsistent probabilistic knowledge bases. The satisfied principles, intended to define the common sense, are founded upon the work in [6] that deals with consistent knowledge base. To our knowledge, our principled tools are the first to tolerate inconsistent probabilistic knowledge bases. We also explain that, even though these tools are usually intractable, they may be employed for consensus decision making. In future research, we should state further or stronger principles to characterise these tools, then investigate tractable approximation of them, in order to obtain a viable and sound methodology for onboard decision making.

REFERENCES

1. W. Hormann and J. Leydold. Automatic random variate generation for simulation input. In *The 2000 Winter Simulation Conference*, pages 675–682. IEEE Press, 2000.

2. Gabriele Kern-Isberner and Thomas Lukasiewicz. Combining probabilistic logic programming with the power of maximum entropy. *Artif. Intell.*, 157(1-2):139–202, 2004.

3. László Lovász and Santosh Vempala. Fast algorithms for logconcave functions: Sampling, rounding, integration and optimization. In *FOCS*. IEEE Press, 2006.

4. László Lovász and Santosh Vempala. The geometry of logconcave functions and sampling algorithms. *Random Struct. Algorithms*, 30(3):307–358, 2007.

5. Cheng Ma, Chao Liu, Shaoxian Ma, and Chengshun Jiang. The application of a new entropy function and mutative scale chaos optimization strategy in two-dimensional entropic image segmentation. In *Computational Intelligence and Security*, volume 2, pages 1647–1652. IEEE Press, Nov. 2006.

6. J.B. Paris. *The Uncertain Reasoner's Companion: A Mathematical Perspective*, volume 39 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1994.

7. J.B. Paris and A. Vencovská. A note on the inevitability of maximum entropy. *International Journal of Approximate Reasoning*, 4(3):183–223, 1990.

8. J.B. Paris and A. Vencovská. In defence of the maximum entropy inference process. *International Journal of Approximate Reasoning*, 17:17–103, 1997.

9. J.B. Paris and A. Vencovská. Common sense and stochastic independence. In Jon Williamson David Corfield, editor, *Fondation of Bayesianism*, number 24 in Applied Logic Series, chapter Logic, Mathematics and Bayesianism, pages 203–240. Kluwer, 2001.

10. Jin Yu, S. V. N. Vishwanathan, Simon Günter, and Nicol N. Schraudolph. A quasi-newton approach to non-smooth convex optimization. In *ICML*, pages 1216–1223, 2008.