An Algebraic Approach to Stochastic ASP

Francisco Coelho [*] Bruno Dinis [†] Salvador Abreu [‡]
fc@uevora.pt bruno.dinis@uevora.pt spa@uevora.pt

October 3, 2023

Abstract

We address the problem of extending probability from the total choices of an ASP program to the stable models and, from there, to general events. Our approach is algebraic in the sense that it relies on an equivalence relation over the set of events and uncertainty is expressed with variables and polynomial expressions. We illustrate our methods with two examples, one of which shows a connection to bayesian networks.

1 Introduction and Motivation

A major limitation of logical representations in real world applications is the implicit assumption that the background knowledge is perfect. This assumption is problematic if data is noisy, which is often the case. Here we aim to explore how answer set programming specifications with probabilistic facts can lead to characterizations of probability functions on the specification's domain, which is not straightforward in the context of answer set programmings, as explained below (see also [3, 11]).

Answer set programming (ASP) [7] is a logic programming paradigm based on the stable model (SM) semantics of normal programs (NPs) that can be implemented using the latest advances in SAT solving technology. Unlike ProLog, ASP is a truly declarative language that supports language constructs such as disjointion in the head of a clause, choice rules, and hard and weak constraints.

[*]Universidade de Évora, NOVALINCS, High Performance Computing Chair
[†]Universidade de Évora, CIMA, CMAFcIO
[‡]Universidade de Évora, NOVALINCS

The distribution semantics (DS) [10, 9] is a key approach to extend logical representations with probabilistic reasoning. Probabilistic facts (PFs) are the most basic DS stochastic primitives and take the form of logical facts, a, labelled with probabilities, p, such as $p :: a$; each PF represents a boolean random variable that is true with probability p and false with probability $\overline{p} = 1 - p$. A (consistent) combination of the PFs defines a total choice (TC) $t = \{p :: a, \ldots\}$ such that (with abuse of notation),

$$\mathrm{P}(T = t) = \prod_{a \in t} p \prod_{a \not\in t} \overline{p}. \tag{1}$$

Our goal is to extend this probability, from TCs, to cover the *specification* domain. We use the term "specification" as set of rules and facts, plain and probabilistic, to decouple it from any computational semantics, implied, at least implicitly, by the term "program". We can foresee two key applications of this extended probability:

1. Support probabilistic reasoning/tasks on the specification domain.
2. Also, given a dataset and a divergence measure, the specification can be scored (by the divergence w.r.t. the *empiric* distribution of the dataset), and weighted or sorted amongst other specifications. These are key ingredients in algorithms searching, for example, optimal specifications of a dataset.

Our idea to extend probabilities from total choices starts with the stance that *a specification describes an observable system*, that *the stable models are all the possible states* of that system and that *observations (i.e. events) are stochastic* — one observation can be sub-complete or super-complete, and might not determine the real state of the system. From here, probabilities must be extended from TCs to SMs and then to any event.

This extension process starts with a critical problem, illustrated by the example in section 2, concerning situations where multiple SMs, ab and ac, result from a single TC, a, but there is not enough information (in the specification) to assign a single probability to each SM. We propose to address this issue by using algebraic variables to describe that lack of information and then estimate the value of those variables from empirical data. This lack of uniqueness is also addressed in [3] along a different approach, using credal sets.

In another related work, [11], epistemic uncertainty (or model uncertainty) is considered as a lack of knowledge about the underlying model, that may be mitigated via further observations. This seems to presuppose a bayesian approach to imperfect knowledge in the sense that having further observations allows to improve/correct the model. Indeed, the approach in that work uses Beta distributions in order to be able to learn the full distribution. This approach seems to be specially fitted to being able to tell when some probability lies beneath some given value. Our approach seems to be similar in spirit, while remaining algebraic in the way that the extension of probabilities is addressed.

The example in section 2 uses the code available in the project's repository[1], developed with the *Julia* programming language [1], and the *Symbolics* [5], and *DataFrames* [2] libraries.

2 A Simple but Fruitful Example

In this section we consider a somewhat simple case, which we call the Simple But Fruitful (SBF) example, that showcases the problem of extending probabilities from total choices to stable models and then to events. As mentioned before, the main issue arises from the lack of information in the specification to assign a single probability to each stable model. This becomes a crucial problem in situations where multiple stable models result from a single total choice. We will come back to this example in section 4.1, after we present our proposal for extending probabilities from total choices to stable models in section 3.

Example 1. *Consider the following specification*

$$\begin{aligned} 0.3 :: a&, \\ b \vee c \leftarrow a&. \end{aligned} \tag{2}$$

This specification has three stable models, $\overline{a}$, ab *and* ac *(see fig. 1). While it is straightforward to assume* $\mathrm{P}(\overline{a}) = 0.7$*, there is no obvious, explicit, way to assign values to* $\mathrm{P}(ab)$ *and* $\mathrm{P}(ac)$*. For instance, we can use a parameter* θ *as in*

$$\begin{aligned} \mathrm{P}(ab) &= 0.3\theta, \\ \mathrm{P}(ac) &= 0.3(1-\theta) \end{aligned}$$

to express our knowledge that ab, ac *are events related in a certain way and, simultaneously, our uncertainty about that relation. This uncertainty can then be addressed with the help of adequate distributions, such as empirical distributions from a dataset. We further develop this case in the examples of section 4.*

[1] https://git.xdi.uevora.pt/fc/sasp

If an ASP specification is intended to describe some system then:

1. With a probability set for the stable models, we want to extend it to all the events of the specification domain.

2. In case where some statistical knowledge is available, for example, in the form of a distribution, we consider it as "external" knowledge about the parameters, that doesn't affect the extension procedure described below.

3. Statistical knowledge can be used to estimate parameters and to "score" the specification.

4. If that specification is only but one of many possible candidates then that score can be used, *e.g.* as fitness, by algorithms searching (optimal) specifications of a dataset of observations.

5. If observations are not consistent with the specification, then we ought to conclude that the specification is wrong and must be changed accordingly.

Currently, we are addressing the problem of extending a probability function (possibly using parameters such as θ), defined on the SMs of a specification, to all the events of that specification. This extension must satisfy the Kolmogorov axioms of probability so that probabilistic reasoning is consistent with the ASP specification and follow our interpretation of stable models as the states of an observable system.

So, as states of a system, we assume that SMs are (statistically) disjoint:

Assumption 1. *Stable models are disjoint events: For any $X \subset \mathcal{S}$*

$$\mathrm{P}(X) = \sum_{s \in X} \mathrm{P}(s) \tag{3}$$

Consider the stable models ab, ac from example 1, that result from the clause $b \vee c \leftarrow a$ and the total choice a. Since we intend to associate each stable model with a state of the system, ab and ac should be *disjoint* events. So $b \vee c$ is interpreted as an *exclusive disjunction* and no further statistical relation between b and c is assumed.

By not making distribution assumptions on the clauses of the specification we can state such properties on the semantics of the specification, as we've done in assumption 1.

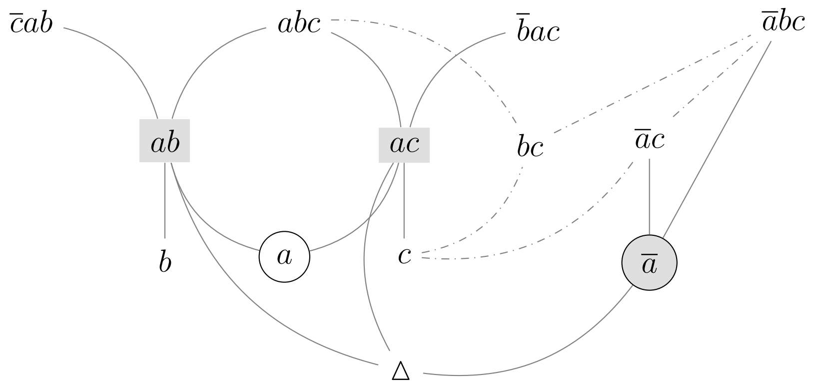

Figure 1: Events related to the stable models of example 1. The circle nodes are total choices and shaded nodes are stable models. The *empty event*, with no literals, is denoted by $\triangle$.

3 Extending Probabilities

The diagram in fig. 1 illustrates the problem of extending probabilities from total choices to stable models and then to general events in a *edge-wise* process, where the value in a node is defined from the values in its neighbors. This quickly leads to coherence problems concerning probability, with no clear systematic approach. Notice that bc is not directly related with any stable model therefore propagating values through edges would assign a hard to justify ($\neq 0$) value to bc. Instead, we propose to base the extension in the relation an event has with the stable models.

3.1 An Equivalence Relation

Given an ASP specification, we consider a set of *atoms* $\mathcal{A}$, a set of *literals*, $\mathcal{L}$, and a set of *events* $\mathcal{E}$ such that $e \in \mathcal{E} \iff e \subseteq \mathcal{L}$. We also consider a set of *worlds* $\mathcal{W}$ (consistent events), a set of *total choices* $\mathcal{T}$ such that for every $a \in \mathcal{A}$ we have $t \in \mathcal{T} \iff t = a \vee \neg a$, and a set of *stable models* $\mathcal{S}$ such that $\mathcal{S} \subset \mathcal{W}$. At last, the set of stable models entailed by the total choice t is denoted by $\langle t \rangle$.

Our path to extend measures and probabilities starts with a perspective of stable models as playing a role similar to *prime factors*. The stable models of a specification are the irreducible events entailed from that specification and any event must be considered under its relation with the stable models.

From example 1, consider the SMs $\{\bar{a}, ab, ac\}$ and events a, abc and c. While a is related (contained) with both ab, ac, event c is related only with

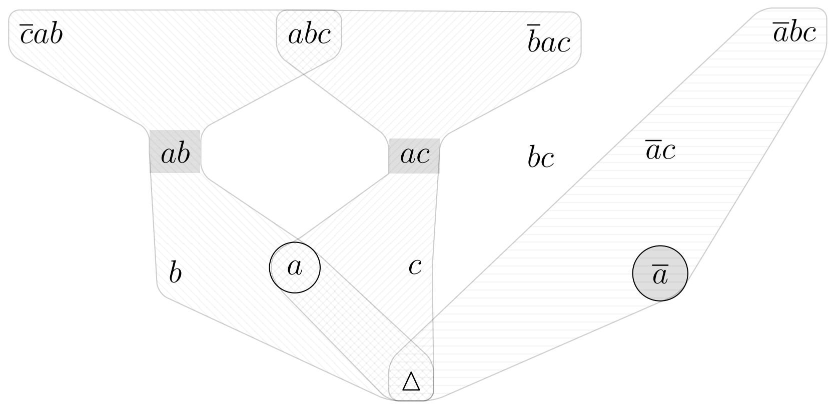

Figure 2: Classes (of consistent events) related to the stable models of example 1 are defined through intersections and inclusions. In this picture we can see, for example, the classes $\{\bar{c}ab, ab, b\}$ and $\{a, abc\}$. As before, the circle nodes are total choices and shaded nodes are stable models. Notice that bc is not in a "shaded" area.

ac. So, a and c are related with different SMs. On the other hand, both ab, ac are related (contained) in abc. So a and abc are related with the same stable models.

Definition 1. *The* stable core (SC) *of the event* $e \in \mathcal{E}$ *is*

$$[\![e]\!] := \{s \in \mathcal{S} \mid s \subseteq e \vee e \subseteq s\}. \tag{4}$$

We now define an equivalence relation, $\sim$, so that two events are related if either both are inconsistent or both are consistent and, in the latter case, with the same stable core.

Definition 2. *For a given specification, let* $u, v \in \mathcal{E}$*. The equivalence relation* $\sim$ *is defined by*

$$u \sim v :\iff u, v \notin \mathcal{W} \vee \left(u, v \in \mathcal{W} \wedge [\![u]\!] = [\![v]\!]\right). \tag{5}$$

Observe that the minimality of stable models implies that, in definition 1, either e is a stable model or at least one of $\exists s\,(s \subseteq e), \exists s\,(e \subseteq s)$ is false. This equivalence relation defines a partition of the events space, where each class holds a unique relation with the stable models. In particular, we denote each class by:

$$[e]_\sim = \begin{cases} \bot := \mathcal{E} \setminus \mathcal{W} & \text{if } e \in \mathcal{E} \setminus \mathcal{W}, \\ \{u \in \mathcal{W} \mid [\![u]\!] = [\![e]\!]\} & \text{if } e \in \mathcal{W}. \end{cases} \tag{6}$$

The combinations of the stable models, together with $\bot$, form a set of representatives. Consider again example 1. As previously mentioned, the stable models are $\mathcal{S} = \{\overline{a}, ab, ac\}$ so the quotient set of this relation, with a small abuse of notation, is:

$$[\mathcal{E}]_{\sim} = \left\{\bot, \Diamond, [\overline{a}]_{\sim}, [ab]_{\sim}, [ac]_{\sim}, [\overline{a}, ab]_{\sim}, [\overline{a}, ac]_{\sim}, [ab, ac]_{\sim}, [\overline{a}, ab, ac]_{\sim}\right\} \tag{7}$$

where $\Diamond$ denotes both the class of *independent* events e such that $[\![e]\!] = \emptyset$ and its core (which is the empty set). We have:

Core, $[\![e]\!]$	**Class**, $[e]_{\sim}$	**Size**, $\#[e]_{\sim}$
$\bot$	$\overline{a}a, \ldots$	37
$\Diamond$	$\overline{b}, \overline{c}, bc, \overline{b}a, \overline{b}c, \overline{bc}, \overline{c}a, \overline{c}b, \overline{bc}a$	9
$\overline{a}$	$\overline{a}, \overline{a}b, \overline{a}c, \overline{ab}, \overline{ac}, \overline{a}bc, \overline{ac}b, \overline{ab}c, \overline{abc}$	9
ab	$b, ab, \overline{c}ab$	3
ac	$c, ac, \overline{b}ac$	3
$\overline{a}, ab$	$\emptyset$	0
$\overline{a}, ac$	$\emptyset$	0
ab, ac	a, abc	2
$\overline{a}, ab, ac$	$\triangle$	1
Ω	all events	64

- Since all events within an equivalence class are in relation with a specific set of stable models, *measures, including probability, should be constant within classes*:
$$\forall u \in [e]_{\sim} \left(\mu(u) = \mu(e)\right).$$

- In general, we have *much more* stable models than literals but their combinations are still *much less* than events. Nevertheless, the equivalence classes allow us to propagate probabilities from total choices to events, as explained in the next subsection.

- In this specific case, instead of dealing with $64 = 2^6$ events, we consider only the $9 = 2^3 + 1$ classes, well defined in terms of combinations of the stable models.

3.2 From Total Choices to Events

Our path to set a distribution on $\mathcal{E}$ starts with the more general problem of extending *measures*, since extending *probabilities* easily follows by means of a suitable normalization (see (15) and (16)), and has two phases:

1. Extension of the probabilities, *as measures*, from the total choices to events.

2. Normalization of the measures on events, recovering a probability.

The "extension" phase, traced by eq. (1) and eqs. (8) to (14), starts with the measure (probability) of total choices, $\mu(t) = \mathrm{P}(T = t)$, expands it to stable models, $\mu(s)$, and then, within the equivalence relation from eq. (5), to (general) events, $\mu(e)$, including (consistent) worlds.

Total Choices. Using eq. (1), this case is given by

$$\mu_{\mathrm{TC}}(t) := \mathrm{P}(T = t) = \prod_{a \in t} p \prod_{a \notin t} \overline{p}. \tag{8}$$

Stable Models. Recall that each total choice t, together with the rules and the other facts of a specification, defines the set $\langle t \rangle$ of stable models associated with that choice. Given a total choice t, a stable model s, and variables or values $\theta_{s,t} \in [0, 1]$, we define

$$\mu(s, t) := \begin{cases} \theta_{s,t} & \text{if } s \in \langle t \rangle \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

such that $\sum_{s \in \langle t \rangle} \theta_{s,t} = 1$.

Classes. Each class is either the inconsistent class, $\bot$, or is represented by some set of stable models.

Inconsistent Class. The inconsistent class contains events that are logically inconsistent, thus should never be observed and have measure zero:

$$\mu(\bot, t) := 0.^{2} \tag{10}$$

Independent Class. A world that neither contains nor is contained in a stable model corresponds to a non-state, according to the specification. So the respective measure is also set to zero:

$$\mu(\Diamond, t) := 0. \tag{11}$$

[2]Notice that this measure being equal to zero is actually independent of the total choice.

Other Classes. The extension must be constant within a class, its value should result from the elements in the stable core, and respects assumption 1 (stable models are disjoint):

$$\mu\big([e]_\sim, t\big) := \mu\big(\llbracket e \rrbracket, t\big) = \sum_{s \in \llbracket e \rrbracket} \mu(s, t) \tag{12}$$

and

$$\mu\big([e]_\sim\big) := \sum_{t \in \mathcal{T}} \mu\big([e]_\sim, t\big)\, \mu_{\mathrm{TC}}(t)\,. \tag{13}$$

Events. Each (general) event e is in the class defined by its stable core, $\llbracket e \rrbracket$. So, denoting by $\#X$ the number of elements in X, we set:

$$\mu(e, t) := \begin{cases} \frac{\mu\left([e]_\sim, t\right)}{\#[e]_\sim} & \text{if } \#[e]_\sim > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{14}$$

and

$$\mu(e) := \sum_{t \in \mathcal{T}} \mu(e, t)\, \mu_{\mathrm{TC}}(t)\,. \tag{15}$$

- Consider the event bc. Since $[bc]_\sim = \Diamond$, from eq. (11) we get $\mu(bc) = 0$.
- Equation (15) together with external statistical knowledge, can be used to learn about the *initial* probabilities of the atoms, that should not (and by proposition 1 can't) be confused with the explicit μ_{TC} set in the specification.

The $\theta_{s,t}$ parameters in equation (9) express the *specification's* lack of knowledge about the measure assignment, when a single total choice entails more than one stable model. In that case, how to distribute the respective measures? Our proposal to address this problem consists in assigning an unknown measure, $\theta_{s,t}$, conditional on the total choice, t, to each stable model s. This approach allows the expression of an unknown quantity and future estimation, given observed data.

Equation (12) results from assumption 1 and states that the measure of a class $[e]_\sim$ is the sum over it's stable core, $\llbracket e \rrbracket$, and (13) *marginalizes* the TCs on (12).

The *normalizing factor* is:

$$Z := \sum_{e \in \mathcal{E}} \mu(e) = \sum_{[e]_\sim \in [\mathcal{E}]_\sim} \mu\big([e]_\sim\big)\,,$$

and now equation (15) provides a straightforward way to define the *probability of observation of a single event*:

$$\mathrm{P}(E = e) := \frac{\mu(e)}{Z}. \tag{16}$$

Since total choices are also events, one can ask, for an arbitrary total choices t, if $\mathrm{P}(T = t) = \mathrm{P}(E = t)$ or, equivalently, if $\mu_{\mathrm{TC}}(t) = \mu(t)$. However, it is easy to see that, in general, that cannot be true. While the domain of the random variable T is the set of total choices, for E the domain is much larger, including all the events. Except for trivial specifications, where the SMs are the TCs, some events other than total choices have non-zero probability.

Proposition 1. *In a specification with a stable model that is not a total choice there is at least one $t \in \mathcal{T}$ such that:*

$$\mathrm{P}(T = t) \neq \mathrm{P}(E = t)\,. \tag{17}$$

Proof. Supposing towards a contradiction that $\mathrm{P}(T = t) = \mathrm{P}(E = t)$ for all $t \in \mathcal{T}$. Then

$$\sum_{t \in \mathcal{T}} \mathrm{P}(E = t) = \sum_{t \in \mathcal{T}} \mathrm{P}(T = t) = 1.$$

Hence $\mathrm{P}(E = x) = 0$ for all $x \in \mathcal{E} \setminus \mathcal{T}$, in contradiction with the fact that for at least one $s \in \mathcal{S} \setminus \mathcal{T}$ one has $\mathrm{P}(E = s) > 0$. □

The essential conclusion of proposition 1 is that we are dealing with *two distributions*: one, on the TCs, explicit in the annotations of the specifications and another one, on the events, and entailed by the explicit annotations *and the structure of the stable models.*

4 Developed Examples

Here we apply the methods from section 3 to the SBF example and to a well known bayesian network, the Earthquake, Burglar, Alarm problem.

4.1 The SBF Example

We continue with the specification from eq. (2).

Total choices. The total choices, and respective stable models, are

Total choice	**Stable models**	$\mu_{\mathbf{TC}}(t)$
a	ab, ac	0.3
$\overline{a}$	$\overline{a}$	$\overline{0.3} = 0.7$

Stable models. The $\theta_{s,t}$ parameters in this example are

$\theta_{s,t}$	$\overline{a}$	a
$\overline{a}$	1	0
ab	0	θ
ac	0	$\overline{\theta}$

with $\theta \in [0, 1]$.

Classes. Following the definitions in eqs. (4) to (6) and (10) to (12) we get the following quotient set (ignoring $\bot$ and $\Diamond$), and measures:

$[\![e]\!]$	$\mu(s, \overline{a})$ $\overline{a}, ab, ac$	$\mu(s, a)$ $\overline{a}, ab, ac$	$\mu\big([e]_\sim, \overline{a}\big)$ $\mu_{\mathrm{TC}} = 0.7$	$\mu\big([e]_\sim, a\big)$ $\mu_{\mathrm{TC}} = 0.3$	$\mu\big([e]_\sim\big)$
$\overline{a}$	$\boxed{1}, 0, 0$	$\boxed{0}, \theta, \overline{\theta}$	1	0	0.7
ab	$1, \boxed{0}, 0$	$0, \boxed{\theta}, \overline{\theta}$	0	θ	0.3θ
ac	$1, 0, \boxed{0}$	$0, \theta, \boxed{\overline{\theta}}$	0	$\overline{\theta}$	$0.3\overline{\theta}$
$\overline{a}, ab$	$\boxed{1}, \boxed{0}, 0$	$\boxed{0}, \boxed{\theta}, \overline{\theta}$	1	θ	$0.7 + 0.3\theta$
$\overline{a}, ac$	$\boxed{1}, 0, \boxed{0}$	$\boxed{0}, \theta, \boxed{\overline{\theta}}$	1	$\overline{\theta}$	$0.7 + 0.3\overline{\theta}$
ab, ac	$1, \boxed{0}, \boxed{0}$	$0, \boxed{\theta}, \boxed{\overline{\theta}}$	0	$\theta + \overline{\theta} = 1$	0.3
$\overline{a}, ab, ac$	$\boxed{1}, \boxed{0}, \boxed{0}$	$\boxed{0}, \boxed{\theta}, \boxed{\overline{\theta}}$	1	$\theta + \overline{\theta} = 1$	1

Prior Distributions. Following the above values (in rational form), and considering the inconsistent and independent classes (resp. $\bot, \Diamond$):

$[\![e]\!]$	$\#[e]_\sim$	$\mu([e]_\sim)$	$\mu(e)$	$\mathrm{P}(E = e)$	$\mathrm{P}(E \in [e]_\sim)$
$\bot$	37	0	0	0	0
$\Diamond$	9	0	0	0	0
$\overline{a}$	9	$\frac{7}{10}$	$\frac{7}{90}$	$\frac{7}{207}$	$\frac{7}{23}$
ab	3	$\frac{3}{10}\theta$	$\frac{1}{10}\theta$	$\frac{1}{23}\theta$	$\frac{3}{23}\theta$
ac	3	$\frac{3}{10}\overline{\theta}$	$\frac{1}{10}\overline{\theta}$	$\frac{1}{23}\overline{\theta}$	$\frac{3}{23}\overline{\theta}$
$\overline{a}, ab$	0	$\frac{7+3\theta}{10}$	0	0	0
$\overline{a}, ac$	0	$\frac{7+3\overline{\theta}}{10}$	0	0	0
ab, ac	2	$\frac{3}{10}$	$\frac{3}{20}$	$\frac{3}{46}$	$\frac{3}{23}$
$\overline{a}, ab, ac$	1	1	1	$\frac{10}{23}$	$\frac{10}{23}$
			$Z = \frac{23}{10}$		

So the prior distributions, denoted by the random variable E, of events and classes are:

$[\![e]\!]$	$\bot$	$\Diamond$	$\overline{a}$	ab	ac	$\overline{a}, ab$	$\overline{a}, ac$	ab, ac	$\overline{a}, ab, ac$
$\mathrm{P}(E = e)$	0	0	$\frac{7}{207}$	$\frac{1}{23}\theta$	$\frac{1}{23}\overline{\theta}$	0	0	$\frac{3}{46}$	$\frac{10}{23}$
$\mathrm{P}(E \in [e]_\sim)$	0	0	$\frac{7}{23}$	$\frac{3}{23}\theta$	$\frac{3}{23}\overline{\theta}$	0	0	$\frac{3}{23}$	$\frac{10}{23}$

(18)

Testing the Prior Distributions

These results can be *tested by simulation* in a two-step process, where (1) a "system" is *simulated*, to gather some "observations" and then (2) empirical distributions from those samples are *related* with the prior distributions from eq. (18). Tables 1 and 2 summarize some of those tests, where datasets of $n = 1000$ observations are generated and analyzed.

Simulating a System. Following some criteria, more or less related to the given specification, a set of events, that represent observations, is generated. Possible simulation procedures include:

- *Random.* Each sample is a Random Set of Literals (RSL). Additional sub-criteria may require, for example, consistent events, a Random Consistent Event (RCE) simulation.
- *Model+Noise.* Gibbs' sampling [4] tries to replicate the specification model and also to add some noise. For example, let $\alpha, \beta, \gamma \in [0, 1]$ be

some parameters to control the sample generation. The first parameter, α is the "out of model" samples ratio; β represents the total choice a or $\overline{a}$ (explicit in the model) and γ is the simulation representation of θ. A single sample is then generated following the probabilistic choices below:

$$\begin{cases} \alpha & \text{by RCE} \\ & \begin{cases} \beta & \overline{a} \\ & \begin{cases} \gamma & ab \\ & ac \end{cases} \end{cases} \end{cases},$$

where

$$\begin{cases} p & x \\ & y \end{cases}$$

denotes "*the value of x with probability p, otherwise y*" — notice that y might entail x and *vice-versa*: E.g. some ab can be generated in the RCE.

- *Other Processes.* Besides the two sample generations procedures above, any other processes and variations can be used. For example, requiring that one of $x, \overline{x}$ literals is always in a sample or using specific distributions to guide the sampling of literals or events.

Relating the Empirical and the Prior Distributions. The data from the simulated observations is used to test the prior distribution. Consider the prior, $\mathrm{P}(E)$, and the empirical, $\mathrm{P}(S)$, distributions and the following error function:

$$\mathrm{err}(\theta) := \sum_{e \in \mathcal{E}} \big(\mathrm{P}(E = e) - \mathrm{P}(S = e)\big)^2. \tag{19}$$

- Since E depends on θ, one can ask how does the error varies with θ.
- What is the *optimal* (i.e. minimum) error value

$$\hat{\theta} := \arg\min_{\theta} \mathrm{err}(\theta) \tag{20}$$

  and what does it tell us about the specification.

In order to illustrate this analysis, consider the experiment summarized in table 1:

1. Equation (19) becomes

$$\mathrm{err}(\theta) = \frac{20869963}{66125000} + \frac{477}{52900}\theta + \frac{18}{529}\theta^2.$$

$[\![e]\!]$	$\#\{S \in [e]_\sim\}$	$\mathrm{P}(S \in [e]_\sim)$	$\mathrm{P}(E \in [e]_\sim)$
$\bot$	0	0	0
$\Diamond$	24	$\frac{24}{1000}$	0
$\overline{a}$	647	$\frac{647}{1000}$	$\frac{7}{23}$
ab	66	$\frac{66}{1000}$	$\frac{3}{23}\theta$
ac	231	$\frac{231}{1000}$	$\frac{3}{23}\overline{\theta}$
$\overline{a}, ab$	0	0	0
$\overline{a}, ac$	0	0	0
ab, ac	7	$\frac{7}{1000}$	$\frac{3}{23}$
$\overline{a}, ab, ac$	25	$\frac{25}{1000}$	$\frac{10}{23}$
	$n = 1000$		

Table 1: **Experiment 1.** Results from an experiment where $n = 1000$ samples where generated following the *Model+Noise* procedure with parameters $\alpha = 0.1, \beta = 0.3, \gamma = 0.2$. The *empirical* distribution is represented by the random variable S while the *prior*, as before, is denoted by E.

2. The minimum of $\mathrm{err}(\theta)$ is at $\frac{477}{52900} + 2\frac{18}{529}\theta = 0$. Since this value is negative and $\theta \in [0, 1]$, it must be $\hat{\theta} = 0$, and

$$\mathrm{err}\left(\hat{\theta}\right) = \frac{20869963}{66125000} \approx 0.31561.$$

The parameters α, β, γ of that experiment favour ac over ab. In particular, setting $\gamma = 0.2$ means that in the simulation process, choices between ab and ac favour ac 4 to 1. For completeness sake, we also describe one experiment that favours ab over ac (setting $\gamma = 0.8$) and one balanced ($\gamma = 0.5$).

For $\gamma = 0.8$**,** the error function is

$$\mathrm{err}(\theta) = \frac{188207311}{529000000} - \frac{21903}{264500}\theta + \frac{18}{529}\theta^2 \approx 0.35579 - 0.08281\theta + 0.03403\theta^2$$

and, with $\theta \in [0, 1]$ the minimum is at $-0.08281 + 0.06805\theta = 0$, *i.e.*:

$$\hat{\theta} : \; \frac{0.08281}{0.06805} \approx 1.21683 > 1. \quad \text{So, } \hat{\theta} = 1,$$
$$\mathrm{err}\left(\hat{\theta}\right) \approx \quad 0.30699.$$

For $\gamma = 0.5$**,** the error function is

$$\mathrm{err}(\theta) = \frac{10217413}{33062500} - \frac{2181}{66125}\theta + \frac{18}{529}\theta^2 \approx 0.30903 - 0.03298\theta + 0.03402\theta^2$$

$[\![e]\!]$	$\#\left\{S_{0.2} \in [e]_\sim\right\}$	$\#\left\{S_{0.8} \in [e]_\sim\right\}$	$\#\left\{S_{0.5} \in [e]_\sim\right\}$
$\perp$	0	0	0
$\Diamond$	24	28	23
$\overline{a}$	647	632	614
ab	66	246	165
ac	231	59	169
$\overline{a}, ab$	0	0	0
$\overline{a}, ac$	0	0	0
ab, ac	7	8	4
$\overline{a}, ab, ac$	25	27	25

Table 2: **Experiments 2 and 3.** Results from experiments where, in each, $n = 1000$ samples are generated following the *Model+Noise* procedure with parameters $\alpha = 0.1, \beta = 0.3, \gamma = 0.8$ (Experiment 2) and $\gamma = 0.5$ (Experiment 3). Empirical distributions are represented by the random variables $S_{0.8}$ and $S_{0.5}$ respectively. Data from experience table 1 is also included, and denoted by $S_{0.2}$, to provide reference.

and, with $\theta \in [0, 1]$ the minimum is at $-0.03298 + 0.068804\theta = 0$, *i.e.*:

$$\begin{aligned}\hat{\theta} &\approx \frac{0.03298}{0.06804} \approx 0.48471 \approx \frac{1}{2}, \\ \operatorname{err}\left(\hat{\theta}\right) &\approx 0.30104\end{aligned}$$

These experiments show that data can indeed be used to estimate the parameters of the model. However, we observe that the estimated $\hat{\theta}$ has a tendency to *exaggerate* the bias of the θ used to generate the samples. More precisely, in experiment 1 data is generated with $\gamma = 0.2$ (the surrogate of θ) and the estimation leads to $\hat{\theta} = 0$ while in experiment 2, $\gamma = 0.8$ leads to $\hat{\theta} = 1$. This suggests that we might need to refine the error estimation process. However, experiment 3 data results from $\gamma = 0.5$ and we've got $\hat{\theta} \approx 0.48471 \approx 0.5$, which is more in line with what is to be expected.

4.2 An Example Involving Bayesian Networks

As it turns out, our framework is suitable to deal with more sophisticated cases, in particular cases involving Bayesian networks. In order to illustrate

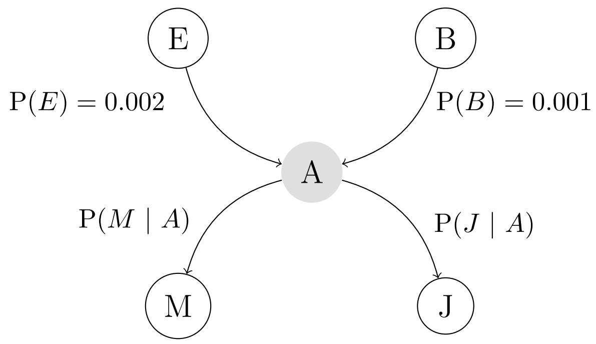

P($M \mid A$)

	m	$\neg m$
a	0.9	0.1
$\neg a$	0.05	0.95

P($J \mid A$)

	j	$\neg j$
a	0.7	0.3
$\neg a$	0.01	0.99

P($A \mid B \wedge E$)

		a	$\neg a$
b	e	0.95	0.05
b	$\neg e$	0.94	0.06
$\neg b$	e	0.29	0.71
$\neg b$	$\neg e$	0.001	0.999

Figure 3: The Earthquake, Burglary, Alarm model

this, in this section we see how the classical example of the Burglary, Earthquake, Alarm [8] works in our setting. This example is a commonly used example in Bayesian networks because it illustrates reasoning under uncertainty. The gist of the example is given in fig. 3. It involves a simple network of events and conditional probabilities.

The events are: Burglary (B), Earthquake (E), Alarm (A), Mary calls (M) and John calls (J). The initial events B and E are assumed to be independent events that occur with probabilities $\mathrm{P}(B)$ and $\mathrm{P}(E)$, respectively. There is an alarm system that can be triggered by either of the initial events B and E. The probability of the alarm going off is a conditional probability given that B and E have occurred. One denotes these probabilities, as per usual, by $\mathrm{P}(A \mid B)$, and $\mathrm{P}(A \mid E)$. There are two neighbors, Mary and John who have agreed to call if they hear the alarm. The probability that they do actually call is also a conditional probability denoted by $\mathrm{P}(M \mid A)$ and $\mathrm{P}(J \mid A)$, respectively.

We follow the convention of representing the (upper case) random variable X by the lower case x. Considering the probabilities given in fig. 3 we obtain the following specification:

$$0.001 :: b,$$
$$0.002 :: e,$$

For the table giving the probability $\mathrm{P}(M \mid A)$ we obtain the specification:

$$0.9 :: p_{m|a},$$
$$0.05 :: p_{m|\overline{a}},$$
$$m \leftarrow a \wedge p_{m|a},$$
$$m \leftarrow \neg a \wedge p_{m|\overline{a}}.$$

This latter specification can be simplified by writing $0.9 :: m \leftarrow a$ and $0.05 :: m \leftarrow \neg a$.

Similarly, for the probability $\mathrm{P}(J \mid A)$ we obtain

$$0.7 :: j \leftarrow a, \quad 0.01 :: j \quad \leftarrow \neg a,$$

Finally, for the probability $\mathrm{P}(A \mid B \wedge E)$ we obtain

$$0.95 :: a \leftarrow b, e, \quad 0.94 :: a \leftarrow b, \overline{e},$$
$$0.29 :: a \leftarrow \overline{b}, e, \quad 0.001 :: a \leftarrow \overline{b}, \overline{e}.$$

One can then proceed as in the previous subsection and analyze this example. The details of such analysis are not given here since they are analogous, albeit admittedly more cumbersome.

5 Discussion and Future Work

This work is a first venture into expressing probability distributions using algebraic expressions derived from a logical specification. We would like to point out that there is still much to explore concerning the full expressive power of logic programs and ASP specifications. So far, we have not considered recursion, logical variables or functional symbols. Also, there is still little effort to articulate with the related fields, probabilistic logical programming, machine learning, inductive programming, *etc.*

The equivalence relation from definition 2 identifies the $s \subseteq e$ and $e \subseteq s$ cases. Relations that distinguish such cases might enable better relations between the models and processes from the stable models.

The example from section 4.2 shows that the theory, methodology, and tools, from bayesian networks can be adapted to our approach. The connection with Markov Fields [6] is left for future work. An example of a "specification selection" application (as mentioned in item 4, section 2) is also left for future work.

Related with the remark at the end of section 4.1 on the tendency of the estimated $\hat{\theta}$ to exaggerate the bias of θ, notice that the error function in (19) expresses only one of many possible "distances" between the empirical and prior distributions. Variations include normalizing this function by the size of $\mathcal{E}$ or using the Kullback-Leibler divergence. The key contribution of this function in this work is to find an optimal θ.

We decided to set the measure of inconsistent events to 0 but, maybe, in some cases, we shouldn't. For example, since observations may be affected by noise, one can expect inconsistencies between the literals of an observation to occur.

Acknowledgements

This work is supported by NOVALINCS (UIDB/04516/2020) with the financial support of FCT.IP. The third author acknowledges the support of FCT - Fundação para a Ciência e Tecnologia under the projects: UIDP/04561/2020 and UIDP/04674/2020, and the research centers CMAFcIO – Centro de Matemática, Aplicações Fundamentais e Investigação Operacional and CIMA – Centro de Investigação em Matemática e Aplicações.

References

[1] Jeff Bezanson et al. "Julia: A Fresh Approach to Numerical Computing". In: *SIAM Review* 59.1 (2017), pp. 65–98. DOI: 10.1137/141000671. eprint: https://doi.org/10.1137/141000671. URL: https://doi.org/10.1137/141000671.

[2] Milan Bouchet-Valat and Bogumił Kamiński. "DataFrames.jl: Flexible and Fast Tabular Data in Julia". In: *Journal of Statistical Software* 107.4 (2023), pp. 1–32. DOI: 10.18637/jss.v107.i04. URL: https://www.jstatsoft.org/index.php/jss/article/view/v107i04.

[3] Fabio Gagliardi Cozman and Denis Deratani Mauá. "The joy of probabilistic answer set programming: semantics, complexity, expressivity, inference". In: *International Journal of Approximate Reasoning* 125 (2020), pp. 218–239.

[4] Stuart Geman and Donald Geman. "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-6.6 (1984), pp. 721–741. DOI: 10.1109/TPAMI.1984.4767596.

[5] Shashi Gowda et al. "High-performance symbolic-numerics via multiple dispatch". In: *arXiv preprint arXiv:2105.03949* (2021).

[6] Ross Kindermann and J. Laurie Snell. *Markov random fields and their applications*. Vol. 1. Contemporary Mathematics. American Mathematical Society, Providence, RI, 1980, pp. ix+142. ISBN: 0-8218-5001-6.

[7] Vladimir Lifschitz. "Answer set programming and plan generation". In: *Artificial Intelligence* 138.1 (2002), pp. 39–54. ISSN: 0004-3702. DOI: https://doi.org/10.1016/S0004-3702(02)00186-8. URL: https://www.sciencedirect.com/science/article/pii/S0004370202001868.

[8] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. The Morgan Kaufmann Series in Representation and Reasoning. Morgan Kaufmann, San Mateo, CA, 1988, pp. xx+552. ISBN: 0-934613-73-7.

[9] Fabrizio Riguzzi. *Foundations of Probabilistic Logic Programming: Languages, Semantics, Inference and Learning*. en. 1st ed. New York: River Publishers, Sept. 2022. ISBN: 978-1-00-333819-2. DOI: 10.1201/9781003338192. URL: https://www.taylorfrancis.com/books/9781003338192 (visited on 03/01/2023).

[10] Taisuke Sato. "A Statistical Learning Method for Logic Programs with Distribution Semantics". In: *International Conference on Logic Programming*. 1995. URL: https://api.semanticscholar.org/CorpusID:10424169.

[11] Victor Verreet et al. "Inference and learning with model uncertainty in probabilistic logic programs". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. 9. 2022, pp. 10060–10069.