International Journal of Approximate Reasoning 69 (2016) 106–132

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

On efficiently estimating the probability of extensions in abstract argumentation frameworks ☆,☆☆

Bettina Fazzinga [a], Sergio Flesca [b,*], Francesco Parisi [b]

[a] ICAR-CNR, Rende (CS), Italy
[b] DIMES – University of Calabria, Rende (CS), Italy

ARTICLE INFO

Article history:
Received 31 July 2015
Received in revised form 19 November 2015
Accepted 23 November 2015
Available online 27 November 2015

Keywords:
Probabilistic reasoning
Argumentation systems
Monte-Carlo simulation

ABSTRACT

Probabilistic abstract argumentation is an extension of Dung's abstract argumentation framework with probability theory. In this setting, we address the problem of computing the probability $Pr^{sem}(S)$ that a set S of arguments is an *extension* according to a semantics *sem*. We focus on four popular semantics (i.e., *complete*, *grounded*, *preferred* and *ideal-set*) for which the state-of-the-art approach is that of estimating $Pr^{sem}(S)$ by using a Monte-Carlo simulation technique, as computing $Pr^{sem}(S)$ has been proved to be intractable. In this paper, we propose a new Monte-Carlo simulation approach which exploits some properties of the above-mentioned semantics for estimating $Pr^{sem}(S)$ using much fewer samples than the state-of-the-art approach, resulting in a significantly more efficient estimation technique.

1. Introduction

Argumentation allows disputes to be modeled, which arise between two or more parties, each of them providing arguments to assert their reasons. The simplest argumentation framework is the *abstract argumentation framework* (AAF) introduced in the seminal paper [2]. An AAF is a pair $\langle A, D \rangle$ consisting of a set A of *arguments*, and of a binary relation D over A, called the *defeat* (or, equivalently, *attack*) relation. Roughly speaking, an argument is an abstract entity representing an assertion of a party, while an attack from an argument a to an argument b indicates that b is contradicted by a (i.e., if a holds then b cannot hold).

Example 1. Consider the following scenario, where we are interested in deciding whether to organize a BBQ party in our garden. Assume that our arguments are the following:

1. a = Our friends will have great fun at the BBQ party;
2. b = Saturday will rain;
3. c = We will not be serving alcoholics at the party.

☆ An abridged version of this paper appeared in [1].

☆☆ The first and the second author were partially supported by the project PON01_01286 – eJRM (electronic Justice Relationship Management).

* Corresponding author.
E-mail addresses: fazzinga@icar.cnr.it (B. Fazzinga), flesca@dimes.unical.it (S. Flesca), fparisi@dimes.unical.it (F. Parisi).

http://dx.doi.org/10.1016/j.ijar.2015.11.009
0888-613X/© 2015 Elsevier Inc. All rights reserved.

This scenario can be modeled by the AAF $\mathcal{A}$, whose set of arguments is $\{a, b, c\}$, and whose defeat relation consists of the defeats $\delta_1 = (b, a)$ and $\delta_2 = (c, a)$, meaning that both in the case of rain and in the case that we do not serve alcoholics our friends could not have much fun. □

Several semantics for AAFs, such as *admissible*, *complete*, *grounded*, *preferred*, and *ideal-set*, have been proposed [2–4] to identify "reasonable" sets of arguments, called *extensions*. Basically, each of these semantics corresponds to some properties which "certify" whether a set of arguments can be profitably used to support a point of view in a discussion. For instance, a set S of arguments is an extension according to the admissible semantics if it has two properties: there is no defeat between arguments in S (it is conflict-free), and every argument (outside S) attacking any argument in S is counterattacked by any argument in S. Intuitively enough, the fact that a set S is an extension according to the admissible semantics means that, using the arguments in S, you do not contradict yourself, and you can rebut to anyone who uses any of the arguments outside S to contradict yours. The fundamental problem of verifying whether a set of arguments is an extension according to one of the above-mentioned semantics has been studied in [5,6].

As a matter of fact, in some context, it may happen that there is a degree of uncertainty about the arguments and attacks used by the parties involved in a dispute. Thus, several proposals have been made to model uncertainty in AAFs, by considering weights, preferences, or probabilities associated with arguments and/or defeats. In this regard, [7–10] have recently extended the original Dung framework in order to achieve probabilistic abstract argumentation frameworks (PrAFs), where the probability theory is exploited to model uncertainty of arguments and defeats. In particular, [8] proposed a PrAF where both arguments and defeats are associated with their (marginal) probabilities.

Example 2. A PrAF $\mathcal{F}$ can be obtained from the AAF $\mathcal{A}$ of Example 1 by considering the arguments a, b, and c, and the defeats δ_1 and δ_2 as probabilistic events, having probabilities $Pr(a) = 0.8$, $Pr(b) = 0.3$, $Pr(c) = 0.5$, $Pr(\delta_1) = 0.7$, and $Pr(\delta_2) = 0.2$. Basically, the values of the probabilities of the events associated with the arguments mean that: we are quite certain that our friends will have great fun at the party; we do not trust a lot the weather forecast service forecasting rain; we are undecided about serving alcoholics; we are quite confident that in the case of rain our friends will not have great fun; and we believe that the lack of alcoholics is very unlikely to make our friend unhappy. □

The issue of how to assign probabilities to arguments and defeats in the PrAF proposed in [8], has been deeply investigated in [11,12], where the *justification* and the *premise* perspectives have been introduced. In this paper, we do not address this issue, but we assume that the probabilities of arguments and defeats are given. We deal with the probabilistic counterpart of the problem of verifying whether a set of arguments is an extension according to a semantics, that is, the problem of *determining the probability* $Pr^{sem}_{\mathcal{F}}(S)$ *that a set* S *of arguments is an extension according to a given semantics sem*. To this end, we consider the PrAF proposed in [8], which is based on the notion of *possible world*. Basically, given a PrAF $\mathcal{F}$, a possible world represents a (deterministic) scenario consisting of some subset of the arguments and defeats in $\mathcal{F}$. Hence, a possible world can be viewed as an AAF containing exactly the arguments and the defeats occurring in the represented scenario. For instance, considering the above-introduced PrAF $\mathcal{F}_{\mathcal{A}}$, the possible world $\langle\{a, b\}, \emptyset\rangle$ is the AAF representing the scenario where only a and b occur, while the possible world $\langle\{a, b, c\}, \{\delta_1, \delta_2\}\rangle$ is the AAF representing the scenario where all the arguments and defeats occur.

In [8] it was shown that a PrAF admits a unique probability distribution over the set of possible worlds, which assigns a probability value to each possible world coherently with the probabilities of arguments and defeats, and allows users to derive probabilistic conclusions from the PrAF. The fact that a PrAF admits a unique probability distribution over the set of possible worlds follows from the *independence assumption*, that is: arguments are viewed as pairwise independent probabilistic events and defeats are viewed as probabilistic events conditioned by the occurrence of the arguments they relate, but independent from any other event. As a matter of fact, the independence assumption is widely used when event correlations are unknown or hard to be derived exactly. For instance, in Example 2, the three events associated with arguments can be reasonably assumed independent from one another, and both the defeats can be deemed independent from one another too.

In this work, we rely on the independence assumption as done in [8], so that, as explained above, a unique probability distribution over the set of possible worlds is defined. Once shown that a PrAF admits a unique probability distribution over the set of possible worlds, the probability $Pr^{sem}_{\mathcal{F}}(S)$ is naturally defined as the sum of the probabilities of the possible worlds where the set S of arguments is an extension according to the semantics *sem*. Unfortunately, as pointed out in [13,14], computing $Pr^{sem}_{\mathcal{F}}(S)$ is intractable (actually, $FP^{\#P}$-complete) for the popular semantics *complete*, *grounded*, *preferred* and *ideal-set*. Indeed, for these semantics, the state-of-the-art approach is that of estimating $Pr^{sem}_{\mathcal{F}}(S)$ by a Monte-Carlo simulation approach, as proposed in [8], since the complexity of exactly computing $Pr^{sem}_{\mathcal{F}}(S)$ is prohibitive.

1.1. Main contributions

In this paper, we propose a new Monte-Carlo-based simulation technique for estimating the probability $Pr^{sem}_{\mathcal{F}}(S)$, where *sem* is one of the following semantics: *complete*, *grounded*, *preferred*, *ideal-set*. In more detail, our strategy relies on the fact that an extension is *complete*, *grounded*, *preferred* or *ideal-set* only if it is both *conflict-free* and *admissible*, and on the fact