

A Random Model for Argumentation Framework: Phase Transitions, Empirical Hardness, and Heuristics

Yong Gao

Department of Computer Science, University of British Columbia Okanagan, Kelowna, Canada
yong.gao@ubc.ca

Abstract

We propose and study, theoretically and empirically, a new random model for the abstract argumentation framework (AF). Our model overcomes some intrinsic difficulties of the only random model of directed graphs in the literature that is relevant to AFs, and makes it possible to study the typical-case complexity of AF instances in terms of threshold behaviours and phase transitions. We proved that the probability for a random AF instance to have a stable/preferred extension goes through a sudden change (from 1 to 0) at the threshold of the parameters of the new model $D(n, p, q)$, satisfying the equation $\frac{4q}{(1+q)^2} = p$. We showed, empirically, that in this new model, there is a clear easy-hard-easy pattern of hardness (for a typical backtracking-style exact solvers) associated with the phase transition. Our empirical studies indicated that instances from the new model at phase transitions are much harder than those from an Erdős-Renyi-style model with equal edge density. In addition to being an analytically tractable models for understanding the interplay between problems structures and effectiveness of (branching) heuristics used in practical argumentation solvers, the model can also be used to generate, in a systematic way, non-trivial AF instances with controlled features to evaluate the performance of other AF solvers.

1 Introduction

An abstract argumentation framework (AF) consists of a set of arguments and a pairwise relationship, known as the attack relationship, defined over the arguments [Dung, 1995]. An abstract argumentation framework can thus be modelled by a directed graph $D(V, A)$, where the set of vertices V correspond to the set of arguments and the set of directed edges A together represent the attack relationship. A directed edge $uv \in A$ indicates that the argument u attacks the argument v (or, v is attacked by u). Throughout this paper, we shall use the words “vertex” and “argument” (“edge” and “attack relation”) interchangeably.

Of particular importance in the theory of abstract argumentation is the notion of an **extension** — a subset of ar-

guments that are “collectively acceptable” [Baroni and Giacomin, 2009]. A preferred extension S is a maximal subset of arguments that is conflict-free (i.e., no attack relation exists between members in S) and is admissible (meaning that any argument attacking a member in S is attacked by some member of S). A stable extension S is a subset of arguments that is conflict-free and attack every argument not in S .

Many extension-related algorithmic problems in AF have been shown to be intractable. For example, the decision problems of the existence of preferred extensions and the existence of stable extensions are both NP-complete [Dunne and Wooldridge, 2009]. In recent years, there have been much interest in identifying tractable classes of AF with special graph-theoretic structures [Dvorák *et al.*, 2012; Dunne, 2007] and in designing and implementing backtracking-style algorithms (enhanced by various heuristic branching rules and data reduction rules) to solve these extension-related problems [Charwat *et al.*, 2015; Walicki and Dyrkolbotn, 2012; Bryant and Krause, 2008].

Another approach to understanding intractable algorithmic problems is to study their “typical-case” complexity. In the literature of artificial intelligence and theoretical computer science, one of the fruitful approaches to typical-case complexity is to study the probabilistic behaviour of random instances drawn from a reasonable probabilistic model [Achlioptas and Peres, 2003; Gomes and Walsh, 2006]. By analysing the threshold phenomena of the solution probability and the associated easy-hard-easy pattern, much insight has been gained on the effectiveness of heuristics that are widely used to tackle problems such as Boolean satisfiability, the constraint satisfaction problem, and graph coloring [Culberson and Gent, 2001; Gomes and Walsh, 2006].

In this regard, there has not been much work on the probabilistic and algorithmic behaviour of random instances of AFs. The only result in the random graph literature that is relevant to abstract argumentation is a theorem by Vega [1990] showing that an Erdős-Renyi-style random directed graph has a kernel with high probability (**whp**¹), as being cited in [Dimopoulos and Torres, 1996; Dunne and Bench-Capon, 2002; Dunne and Wooldridge, 2009]. Dunne and Wooldridge [2009, Section 6.1] discussed the need to

¹By **whp**, we mean “with probability of the event approaching one as the number of vertices tends to infinity.”

study the average-case properties of an argumentation framework and called for an investigation into the implication of Vega’s result which, when translated to the argumentation context, states that a random AF, where each of the possible attack relationship appears independently with a constant probability, almost always has a stable extension. More recently, Dunne and Atkinson [2014] demonstrated (empirically) a threshold behavior of a random model for the value-based argumentation framework and Wen *et al.* [2016] studied a random model for a special class of answer-set programs where self-loops plays an important role ².

In this work, we have successfully identified a crucial characteristics of randomly-generated AFs and proposed a new random model of AFs with a theoretically-proven phase transition of the solution probability and empirically-verified easy-hard-easy pattern about the complexity of finding an extension associated with the phase transition.

We proved theoretically the exact threshold of our model parameters for the phase transition, which also provides a plausible interpretation of Vega’s result — what makes Vega’s model almost always has a stable extension is that on average, the model produces the “right” fraction of symmetric relationships regardless of the probability of an attack relation.

Empirically, we designed and implemented an effective backtracking-style exact algorithm together with several non-trivial branching rules and heuristics, which are of interest on their own. Using our implementation, we observed a clear phase transition of the solution probability of our model and an easy-hard-easy pattern of hardness of random instances from our model associated with the phase transition. Our experiments also showed that random instances from our model at the phase transitions are typically much harder than those from the random directed graph studied by Vega [1990].

The rest of the paper is organized as follows. In Section 2, we introduce our random model and discuss the main theorem and its consequence on the exact threshold of the phase transition. In Section 3, we analyse the probabilistic behaviour of the random model and prove the main theorem. In Section 4, we report the observations and insights from our empirical studies, including a discussion of the algorithm and heuristics we designed and implemented and the motivation of doing so. We conclude in Section 5 with a discussion on the significance of the work and future research directions.

2 The Model and the Threshold Phenomena

Vega [1990] studied the existence of kernels in a random model $\tilde{D}(n, \tilde{p})$ of directed graphs on n vertices where each of the $n(n-1)$ possible directed edges is present independently with probability \tilde{p} , and showed that for any constant $\tilde{p} > 0$, the probability for $\tilde{D}(n, \tilde{p})$ to have a kernel approaches 1 as n tends to infinity. The result easily translates to the existence of a stable extension in AFs when $\tilde{D}(n, \tilde{p})$ is regarded as a randomly-generated AF [Dimopoulos and Torres, 1996; Dunne and Bench-Capon, 2002; Dunne and Wooldridge, 2009], indicating that for any constant $\tilde{p} > 0$, $\tilde{D}(n, \tilde{p})$ has a stable extension with probability asymptotic to one.

²Thanks to two referees for pointing out these two references.

We propose the following random model $D(n, p, q)$ for AFs with two parameters that explicitly control the likelihood of creating a mutual attack relation between a pair of arguments. As we shall show in this paper, unlike the model studied by Vega, the probability for $D(n, p, q)$ to have a stable/preferred extension (or simply, an admissible set) goes through a sudden change (from 1 to 0) as the model parameter q (conditional probability of a symmetric attack relation) crosses a threshold $q^* = q^*(p)$.

Definition 2.1. $D(n, p, q)$ is defined to be a random abstract argumentation system on n arguments where the $\binom{n}{2}$ pairs of arguments independently decide the existence and the type of the attack relation. Between each pair, an attack relation exists with probability p . Conditional on the existence of an attack relation, the attack relation is a symmetric one (mutual attack) with probability q ; with probability $1 - q$, the attack is a one-way attack and each of the two possible directions are equally likely.

We have the following theorem, characterizing the exact threshold of the phase transition of the probability for the above model to have a non-empty preferred extension.

Theorem 1. For any constants $0 < p < 1$ and $0 < q < 1$,

$$\lim_{n \rightarrow \infty} \mathbb{P}(D(n, p, q) \text{ has a non-empty preferred extension}) = \begin{cases} 1, & \text{if } q > q^* \\ 0, & \text{if } q < q^* \end{cases} \quad (1)$$

where $q^* = q^*(p)$ is the unique value of q between 0 and 1 such that $\frac{4q}{(1+q)^2} = p$.

As the proofs in the next section indicate, the above result also holds for the existence of stable extensions or simply, non-empty admissible sets. It turns out, somewhat surprisingly, that the average number of edges participating in a symmetric attack is closely related to the phase transition of solution probability. Let E_{sym} (or \tilde{E}_{sym}) be the average number of edges in $D(n, p, q)$ (respectively, $\tilde{D}(n, \tilde{p})$) that participate in a symmetric attack relation. By the definition of the two models, we have

$$\frac{E_{sym}}{\tilde{E}_{sym}} = \frac{2 \binom{n}{2} pq}{2 \binom{n}{2} (\tilde{p})^2} = \frac{pq}{(\tilde{p})^2}.$$

To see the impact of the number of symmetric attacks on the solution probability of the two models, we consider the situation where the two models have the same average number of edges (implying that $\tilde{p} = \frac{p(1+q)}{2}$). We have

$$\frac{E_{sym}}{\tilde{E}_{sym}} = \frac{4q}{p(1+q)^2} \begin{cases} > 1, & \text{if } q > q^* \\ = 1, & \text{if } q = q^* \\ < 1, & \text{if } q < q^* \end{cases} \quad (2)$$

This suggests a plausible answer to Dunne and Wooldridge’s question regarding why stable extensions always exist in Vega’s model (Dunne and Wooldridge [2009]) — Vega’s model almost always has a stable extension because on average, the model produces the “right” fraction of symmetric relationships regardless of the probability of an attack relationship.

3 Probabilistic Analysis

In this section, we analyse the probability for the existence of a stable extension, a preferred extension, and an admissible set in $D(n, p, q)$. The results together prove Theorem 1.

3.1 Notation

For notation related to (directed) graphs, we follow Bang-Jensen and Gutin [2008]. The cardinality of a set S is denoted by $|S|$. Given a directed graph $D(V, A)$, the out-neighborhood of a vertex $v \in V$ in a vertex subset S , denoted by $N_S^+(v)$, is the set of arguments in S that are attacked by v . Similarly, the in-neighborhood of v in S , $N_S^-(v)$, is the subset of arguments in S that attack v . A subset S of arguments in V is admissible if it is

- **conflict-free** (i.e. $N_S^-(v) \cup N_S^+(v) = \emptyset, \forall v \in S$), and
- every argument $v \in S$ is **acceptable** with respect to S (i.e., $N_V^-(v) \subseteq N_V^+(S)$).

A maximal admissible set is known as a preferred extension. A subset S of arguments is a stable extension if it is conflict-free and every argument v outside of S is attacked by some argument in S , i.e. $v \in N_V^+(S), \forall v \in V \setminus S$. While the empty set is an extension by definition, all the results in this paper are about the existence of non-empty extensions.

3.2 Stable Extensions Above the Threshold

We show that above the threshold, $D(n, p, q)$ has a stable extension **whp**.

Theorem 2. *For any $0 \leq p, q < 1$, if $q > q^*$, then*

$$\lim_{n \rightarrow \infty} \mathbb{P}(D(n, p, q) \text{ has a stable extension}) = 1. \quad (3)$$

Proof. Let $N_{\text{stable}}(r)$ be the number of size- r stable extensions in $D(n, p, q)$. We have

$$\mathbb{E}[N_{\text{stable}}(r)] = \binom{n}{r} (1-p)^{\binom{r}{2}} \left(1 - \left(1 - \frac{1+q}{2}p\right)^r\right)^{n-r} \quad (4)$$

where the term $1 - \left(1 - \frac{1+q}{2}p\right)^r$ is the probability for a particular argument to be attacked by at least one argument from a given subset of r arguments. The idea is to show that **whp**, $N_{\text{stable}}(r)$ tends to infinity for a particular value of r .

It can be verified that if $q > q^*$, then $\left(1 - \frac{1+q}{2}p\right)^2 < 1 - p$. Therefore, there exists a constant $c < 2$ (with $|2 - c|$ sufficiently small) such that $\left(1 - \frac{1+q}{2}p\right)^c < 1 - p$. First, consider the expected number of stable extensions of size $r = c \log_b n$ where $b = \frac{1}{1-p}$. Based on Equation (4) and the well-known approximation $\binom{n}{r} \sim \frac{1}{\sqrt{2\pi r}} \left(\frac{en}{r}\right)^r$, we see that

$$\begin{aligned} & \mathbb{E}[N_{\text{stable}}(r)] \\ & \sim \frac{1}{\sqrt{2\pi r}} \left(\frac{en}{r}\right)^r \left(\frac{1}{b^{\frac{r-1}{2}}}\right)^r \left(1 - \left(1 - \frac{1+q}{2}p\right)^r\right)^{n-r} \\ & \sim \frac{1}{\sqrt{2\pi r}} \left(\sqrt{ber}^{-1} n^{1-\frac{c}{2}}\right)^{c \log_b n} e^{-n^{1+c \log_b \left(1 - \frac{1+q}{2}p\right)}}, \end{aligned}$$

which tends to infinity since $c < 2$ and $c \log_b \left(1 - \frac{1+q}{2}p\right) < -1$. Then, we use Chebyshev's inequality to show that

$N_{\text{stable}}(r)$ is close to its expectation asymptotically with probability 1. To this end, let I_S be the indicator function of the event that $S \subset V$ is a stable extension. The number of stable extensions $N_{\text{stable}}(r)$ can be written as $N_{\text{stable}}(r) = \sum_S I_S$,

where the summation is over all size- r subsets of arguments. By Chebyshev's inequality, we have

$$\mathbb{P}(N_{\text{stable}}(r) = 0) \leq \frac{\mathbb{E}[N_{\text{stable}}^2(r)]}{(\mathbb{E}[N_{\text{stable}}(r)])^2} - 1.$$

Grouping pairs of size- r vertex sets by the number of vertices they have in common, we get the following expression of $\mathbb{E}[N_{\text{stable}}^2(r)]$.

$$\begin{aligned} \mathbb{E}[N_{\text{stable}}^2(r)] &= \sum_{S, T} \mathbb{E}[I_S I_T] \\ &= \sum_{l=0}^r \binom{n}{r} \binom{r}{l} \binom{n-r}{r-l} (1-p)^{2\binom{r}{2} - \binom{l}{2}} P(l) \\ &= \sum_{l=0}^r R(l), \end{aligned}$$

where $P(l)$ is the probability that two independent sets with l common members are both admissible. We can show that for any $r \leq c \log_b n$ and $q > q^*$,

$$\lim_n \frac{R(0)}{(\mathbb{E}[N_{\text{stable}}(r)])^2} \leq 1$$

and

$$\lim_n \frac{\sum_{l=1}^r R(l)}{(\mathbb{E}[N_{\text{stable}}(r)])^2} = 0.$$

□

3.3 Preferred Extensions Below the Threshold

In this subsection, we prove that below the threshold, the probability for $D(n, p, q)$ to have an admissible set goes to zero as n goes to infinity. The strategy is to show that the expected number of admissible sets goes to zero as n tends to infinity and then, to apply Markov's inequality (see, e.g., [Alon and Spencer, 2000]) to show that the probability goes to zero as well.

Let $N_{\text{adm}}(r)$ be the number of size- r admissible sets. We have $\mathbb{E}[N_{\text{adm}}(r)] = \binom{n}{r} P_{\text{adm}}(r)$, where $P_{\text{adm}}(r)$ is the probability that a particular subset of r arguments is admissible.

Lemma 3.1.

$$P_{\text{adm}}(r) = (1-p)^{\binom{r}{2}} \left(1 - \left(1 - \frac{1+q}{2}p\right)^r + (1-p)^r\right)^{n-r}$$

Proof. Note that a particular subset A of r arguments is an admissible set if and only if it is conflict-free and $\forall v \notin A$,

- either v is attacked by A (i.e., $v \in N^+(A)$) or
- v does not attack any member of A ($N^+(v) \cap A = \emptyset$).

Therefore, the probability for A to be an admissible set is

$$(1-p)^{\binom{r}{2}} (P_{\text{trouble-free}}(r))^{n-r}$$

where $P_{\text{trouble-free}}(r)$ is the probability for an argument $v \notin A$ to be safe in the following sense: v does not attack any member in A , or at least one member in A attacks v . By the definition of the model $D(n, p, q)$, we have

$$P_{\text{trouble-free}}(r) = (1-p)^r + \sum_{k=1}^r \binom{r}{k} p^k (1-p)^{r-k} \left(1 - \left(\frac{1}{2}(1-q)\right)^k\right),$$

where on the right hand side of the above equation, the first term is the probability that no attacking relation exists between v and any member of A , and the second term is the probability that when there are exactly k attack relations between v and A , at least one of the argument in A attacks v . Rearranging the right-hand side and using the fact that $\sum_{k=0}^r \binom{r}{k} p^k (1-p)^{r-k} = 1$, we see that

$$\begin{aligned} P_{\text{trouble-free}}(r) &= 1 - \sum_{k=1}^r \binom{r}{k} \left(\frac{1}{2}p(1-q)\right)^k (1-p)^{r-k} \\ &= 1 - \left(1 - \frac{1+q}{2}p\right)^r + (1-p)^r \end{aligned}$$

□

Theorem 3. *If $q < q^*$, then*

$$\lim_{n \rightarrow \infty} \mathbb{P}(D(n, p, q) \text{ has an admissible set}) = 0.$$

Proof. Let E be the event that there exists an admissible set. By Markov's inequality, we have

$$\mathbb{P}(E) \leq \sum_{r=1}^{2 \log_b n} \mathbb{E}[N_{\text{adm}}(r)] + \sum_{r=1+2 \log_b n}^n \mathbb{E}[N_{\text{adm}}(r)],$$

where we write $b = \frac{1}{1-p}$. Using Lemma 3.1, it can be shown that the second summation on the right-hand-side of the above equation goes to zero for any p and q , while the first summation goes to zero under the condition of the theorem that $q < q^*$; in this case, we have $p > \frac{4q}{(1+q)^2}$, and therefore $1 + 2 \log_b \left(1 - \frac{1+q}{2}p\right) > 0$. □

Proof of Theorem 1. Since a stable extension is a preferred extension and a preferred extension is an admissible set, Theorem 1 follows from Theorem 2 and Theorem 3. □

4 Empirical Studies and the Algorithm

To empirically study the probabilistic behaviour of the model and the associated/anticipated easy-hard-easy patterns of the random instances, we designed and implemented a branch-and-reduce algorithm — backtracking search enhanced by a set of branch/data-reduction rules and branching heuristics. In addition to their applications in practice, algorithms in the branch-reduce paradigm are also intensively analysed

in the theoretical computer science literature [Fomin and Kratsch, 2010]. Widely used in practical and industrial-strength solvers, the performance of this type of algorithms depends heavily on the choice of branch rules, reduction rules, and the branching heuristics.

In the algorithm described in Algorithm 1, we use the following definitions:

- E (Partial Extension) — the extension to be constructed;
- F (Friendly Set) — arguments that do not attack and are not attacked by any argument in E ; these arguments can potentially be added to E ; and
- H (Hostile Set) — arguments that attack arguments in E , but are not attacked by E . These arguments have to be taken care of to extend E to an extension.

Algorithm 1: ArgExtension(E, F, H)

```

if  $H = \emptyset$ , then return  $E$ ;           %  $E$  is admissible
(*)  $P \leftarrow$  select-pivot-set( $F, H$ );
while  $P \neq \emptyset$  do
(**)  $v \leftarrow$  select-pivot( $P$ );
     $E_{\text{new}} \leftarrow E \cup \{v\}$ ;
     $F_{\text{new}} \leftarrow F \setminus (N_F(v) \cup \{v\})$ ;
     $H_{\text{new}} \leftarrow (H \cup (N_F^-(v) \setminus N_F^+(v))) \setminus N_H^+(v)$ ;
     $A \leftarrow$  ArgExtension( $E_{\text{new}}, F_{\text{new}}, H_{\text{new}}$ );
    if  $A \neq \emptyset$ , return  $A$ ;
     $P \leftarrow P \setminus \{v\}$ ;
end
Return  $\emptyset$ ;
    
```

As long as the hostile set is not empty, the algorithm finds a subset $P \subseteq F$ of friendly arguments such that one of the arguments in P has to be in the extension and branches on the arguments in P . It is well-accepted in the filed of heuristic search that the time complexity of such algorithms can be upper bounded by considering its search tree. In this case, each tree node is an argument the algorithm attempts to include in E , and the children of a search-tree node correspond to the nodes in P . The running time of the algorithm is $O(t_n T)$ where T is the size of the search tree and t_n is the time for processing a tree node; in our implementation, $t_n = O(n)$.

Invoking the algorithm with different inputs solves different extension problems: Stable Extension ($E = \emptyset, F = V, H = V$) and Admissible Set Containing a Particular Argument v ($E = \{v\}, F = V \setminus N(v), H = N^-(v)$). Line (*) and line (**) in Algorithm 1 are where different branching rules and heuristics are incorporated:

- BR_{mcha} (most-constrained-hostile-argument branching rule) Select $P = N_F^-(h)$ such that $|N_F^-(h)|$ is minimized over all hostile arguments $h \in H$.
- BH^0 (fixed-order-branching), BH^1 (random-order branching), BH^2 (maximum-out-degree branching), and
- BH^3 (most-effective-argument-first heuristics). Order the arguments v in P by $|N_H^+(v)| - |N_F^-(v) \setminus N_F^+(v)|$.

We note that modern complete SAT solvers, which are highly competitive in solving CNF-encoded argumentation

problems [Thimm *et al.*, 2016], are also branch-and-reduce style algorithms. However, we emphasize that the purpose of our empirical studies is to address and explore questions related to the probabilistic and algorithmic behaviour of the model and their impact on the effectiveness of heuristics that theoretical analysis is not sufficient to answer, rather than simply demonstrating the superior of one solver over the others. Observations and insights obtained in such studies shall benefit the development of all solvers. Because of this, we designed and implemented our own solver due to the following considerations: (1) we want a complete control in our experiments over the use of branching rules and heuristics; and (2) many of these branching rules/heuristics are designed using unique features of the extension-related problems and thus might not be available in a typical SAT-solver-based (or other backtracking-based) AF solvers.

While the algorithm can be customized to solve different argumentation problems, in our empirical studies we focused on the most basic problem of deciding the existence of admissible sets. Experiments were conducted on instances generated from $D(n, p, q)$ with different combinations of the model parameters: $n = 128, 256, 512, 1024, 2048$, $p = 0.181, 0.331, 0.64, 0.75$, and q . The four values for p , with the corresponding theoretical value of the threshold q^* being $\frac{1}{20}, \frac{1}{10}, \frac{1}{4}, \frac{1}{3}$, are chosen to represent both dense and relatively sparse random graphs. The implementation is in Java and experiments were run on desktop computers with an Intel dual-core processor. For each combination of the model parameters, we used our solver to solve 100 instances from $D(n, p, q)$. In the rest of this section, we report our experiments and discuss our observations.

4.1 The Easy-Hard-Easy Pattern

In our experiments with all the parameter values p , we observed clearly a sudden change of the fraction of yes-instances (those that have an admissible set) out of 100 randomly-generated AFs from $D(n, p, q)$ at the value of q close to the theoretical threshold.

More importantly, the anticipated easy-hard-easy pattern of hardness is obvious for all the parameter values. The result for $n = 1024$ and $p = 0.75$ is presented in Figure 1, where the fraction of yes-instances are plotted as a function of q . Also plotted are the median and average of the size of the search tree, showing that there is indeed an easy-hard-easy pattern similar to those observed when using this class of algorithms to solve other algorithmic problems such as SAT, constraint satisfaction problems, and graph colouring.

We also observed that while yes-instances become increasingly harder as q approaches the threshold of the phase transition, no-instances are consistently much harder, with those that show up at the onset of the phase transition being the hardest. This explains the large disparity between the median and the average at $q = 0.22$ and 0.24 ; in these two cases, there are more no-instances than yes-instances (and thus, the median value is determined by the search tree size of some no-instances), but there are still significant fraction of yes-instances which result in a lower average search tree size.

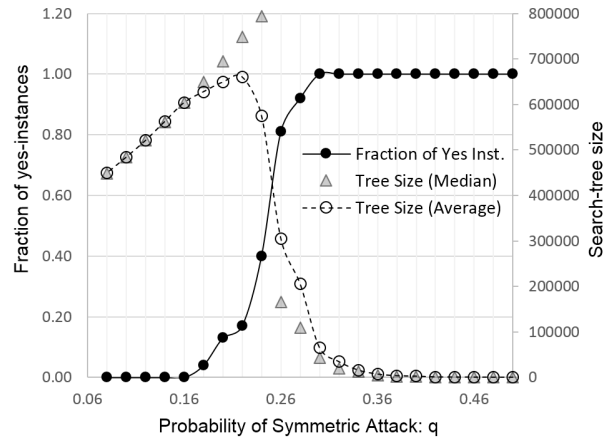


Figure 1: Phase transitions and hardness pattern of $D(n, p, q)$: $n = 1024, p = 0.75$.

4.2 Comparisons of $D(n, p, q)$ and $\tilde{D}(n, \tilde{p})$

We designed and performed two sets of experiments to compare the hardness of randomly-generated instances from our model $D(n, p, q)$ and Vega’s model $\tilde{D}(n, \tilde{p})$,

To compare objectively the typical-case hardness of the two models, several factors have to be taken into consideration, including the number of attack relations, the number of symmetric attack relations, and the size of a possible admissible sets. To guarantee that instances from the two models have the same average number of attack relations, we require that $\tilde{p} = p(1 + q)/2$. In fact, our analyses in Section 3 and those in Vega [1990], can be used to show that the maximum possible sizes of an admissible set are equal in the two models at the phase transition. Therefore, in our experiments, we compare the hardness of random instances from the two models where p and q are in the area of the phase transition with the corresponding $\tilde{p} = p(1 + q)/2$.

In the first set of experiments, we compare the size of the search tree of our solver applied to randomly-generated instances from the two models with $p = 0.75, q = 0.26$, and $n = 128, 256, \dots, 2048$. The selection of these parameters to compare is based on a preliminary experiment with $q = 0.26$ being the observed value where the phase transition starts for $p = 0.75$, and that it is only practically feasible to perform the complete set of experiments for n up to 2048 when p is as large as 0.75.

The results are presented in Figure 2, where we plot the median and the average of the search-tree size as a function of n to solve random AF instances from $D(n, p, q)$ and $\tilde{D}(n, \tilde{p})$, with $p = 0.75, q = 0.26$, and $\tilde{p} = 0.4725$. As the figure shows, the size of the search tree for instances from $D(n, p, q)$ increases at a much higher rate as n increases than those from $\tilde{D}(n, \tilde{p})$. For $n = 2048$, while our implementation solved 100 instances from Vega’s model in a reasonable amount of time (36 hours), instances from $D(n, p, q)$ are much harder — the solver takes more than one week to complete 100 instances at an average speed of about 90 minutes per Yes-instance and more than 4 hours per No-instance.

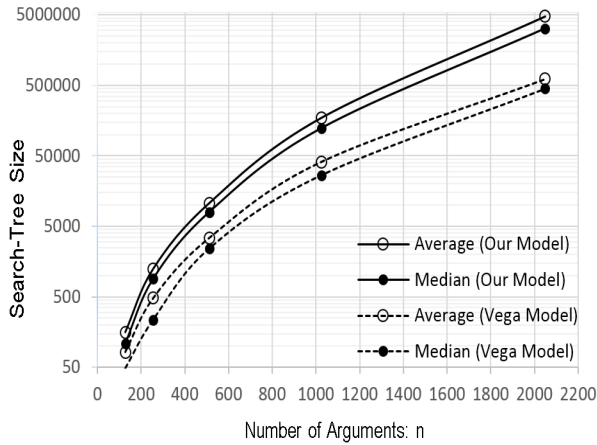


Figure 2: Comparison between our model $D(n, p, q)$ and the Vega model $\tilde{D}(\tilde{p})$: $p = 0.75$, $q = 0.26$, $\tilde{p} = \frac{p(1+q)}{2}$.

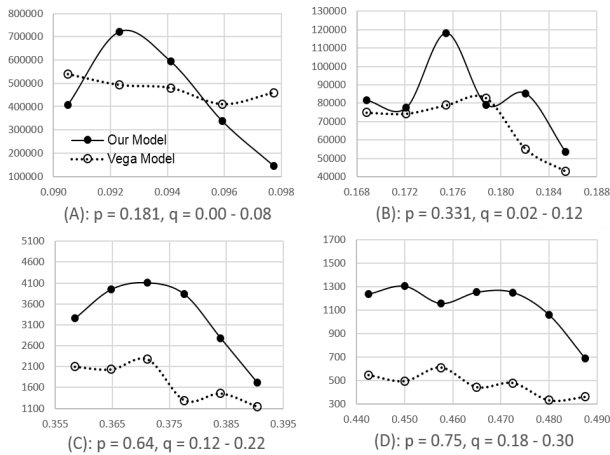


Figure 3: Comparison of our model $D(n, p, q)$ and Vega model $\tilde{D}(\tilde{p})$ with $n = 256$, different \tilde{p} (the horizontal axis), and the corresponding values of p and q around the observed phase transitions. The vertical axis is for the search tree size. Legends apply to all the four plots.

In the second set of experiments, we compare the hardness of the two models with different \tilde{p} of Vega’s model. When \tilde{p} is small (and thus, p must be small due to the equal-density requirement), instances from both models with $n = 512$ are out of reach of our solver on a standard desktop computer. We therefore fixed $n = 256$ and managed to complete the experiments in several days for representative values of \tilde{p} that we are interested in. As is shown in the four plots in Figure 3, the average hardness of instances from the Vega model steadily increases as \tilde{p} decreases. On the other hand, the results shown in the plots indicate that for any $0 < \tilde{p} < 1$, there are corresponding sets of values of p and q such that $D(n, p, q)$ and $\tilde{D}(n, \tilde{p})$ have the same average number of attack relations, but instances from $D(n, p, q)$ are much harder to solve.

4.3 The Importance of Branching Heuristics

Our experiments indicate that the most-constrained-hostile-argument-first branch rule BR is critical. More interestingly, we observed that while the branching heuristic BH^3 improves the performance significantly when the branching rule BR is not used, it is only notably helpful (when BR is used by the solver) before the phase transition and up to the point where hardness of a typical Yes-instance starts to increase dramatically. Due to space limitation, we omit the details of this set of experiments here, but plan to conduct in the future a larger scale and comprehensive empirical study using the proposed model on our solver and other start-of-the-art publicly-available argumentation solvers.

5 Conclusions and Future Work

In an effort to answering the call by Dunne and Wooldridge [2009] to study the average-case behaviour of abstract argumentation and understand the implication of Vega’s result that the Erdős-Renyi-style random model asymptotically almost always has a stable extension, we proposed the random model $D(n, p, q)$ which, to the best of my knowledge, is the first non-trivial random model for the abstract argumentation framework with a proven threshold phenomenon and an empirically-observed easy-hard-easy pattern at the phase transition.

As an immediate future work, it is interesting to conduct a systematic study, using the proposed random model on other state-of-the-art solvers for AFs to gain further insight on the pros and cons of different branching rules and heuristics used in these systems. We emphasize that the main focus of such studies will not be to “demonstrate” which solvers are better than the others; instead, the objective would be to understand the interplay between structural properties and the effectiveness of the branching heuristics and rules, so that insights obtained in such studies can benefit the community at large.

Unlike solution concepts well-studied in the phase-transition literature such as those for SAT and colouring, the existence of extensions in AFs is not monotone in terms of the set of attacks: the property is not preserved by adding more attack relationships. Our random model exploits the fact that the property of having an extension is monotone in the set of symmetric attacks, whereas the random model for answer-set programs by Wen *et al.* [2016] depends heavily on the use of self-contradiction rules and the model for value-based AFs by Dunne and Atkinson [2014] makes use of a special acyclic structures introduced by some transitive order on the set of the values associate with the arguments. A nature question is if there are other interesting features that can be used to introduce different types of random AF models.

In the long term, another future direction of more theoretical flavor is to analyse the typical complexity of random instances of our model, and to explore the possibility of establishing (exponential) lower bounds on the time complexity of backtracking-style AF solvers. See, e.g., [Nordström, 2015], for an informal but up-to-date account of relevant results in proof complexity and the insights of these results may provide for developing practical solvers for various inference systems.

References

- [Achlioptas and Peres, 2003] D. Achlioptas and Y. Peres. The threshold for random k -SAT is $2k(\ln 2 - o(k))$. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC'03)*, pages 223–231, 2003.
- [Alon and Spencer, 2000] N. Alon and J. Spencer. *The Probabilistic Method*. Wiley, 2000.
- [Bang-Jensen and Gutin, 2008] J. Bang-Jensen and G. Gutin. *Digraphs: Theory, Algorithms and Applications*. Springer, 2008.
- [Baroni and Giacomin, 2009] P. Baroni and M. Giacomin. Semantics of abstract argument systems. In I. Rahwan and G. Simari, editors, *Argumentation in Artificial Intelligence*, pages 25–44. Springer, 2009.
- [Bryant and Krause, 2008] D. Bryant and P. Krause. A review of current defeasible reasoning implementation. *The Knowledge Engineering Review*, 23(3):227–260, 2008.
- [Charwat *et al.*, 2015] G. Charwat, W. Dvorak, S. Gaggl, J. Wallner, and S. Woltran. Methods for solving reasoning problems in abstract argumentation - a survey. *Artificial Intelligence*, 220:28–63, 2015.
- [Culberson and Gent, 2001] J. Culberson and I. Gent. Frozen development in graph coloring. *Theoretical Computer Science*, 265(1-2):227–264, 2001.
- [Dimopoulos and Torres, 1996] Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170:209–244, 1996.
- [Dung, 1995] P. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n -person games. *Artificial Intelligence*, 77:321–357, 1995.
- [Dunne and Atkinson, 2014] Paul E. Dunne and K. Atkinson. Properties of random vafs and implications for efficient algorithms. In *Proceedings of Computational Models of Argument (COMMA'14)*, pages 261–268, 2014.
- [Dunne and Bench-Capon, 2002] P. Dunne and T. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141:187–203, 2002.
- [Dunne and Wooldridge, 2009] P. Dunne and M. Wooldridge. Complexity of abstract argumentation. In I. Rahwan and G. Simari, editors, *Argumentation in Artificial Intelligence*, pages 85–103. Springer, 2009.
- [Dunne, 2007] P. Dunne. Computational properties of argument systems satisfying graph-theoretical constraints. *Artificial Intelligence*, 171:701–729, 2007.
- [Dvorák *et al.*, 2012] W. Dvorák, S. Ordyniak, and S. Szeider. Augmenting tractable fragments of abstract argumentation. *Artificial Intelligence*, 186:157–173, 2012.
- [Fomin and Kratsch, 2010] F. Fomin and D. Kratsch. *Exact Exponential Algorithms*. Springer, 2010.
- [Gomes and Walsh, 2006] C. Gomes and T. Walsh. Randomness and structure. In F. Rossi, P. van Beek, and T. Walsh, editors, *Handbook of Constraint Programming*, pages 639–664. Elsevier, 2006.
- [Nordström, 2015] J. Nordström. On the interplay between proof complexity and SAT solving. *ACM SIGLOG News*, 2(3), 2015.
- [Thimm *et al.*, 2016] M. Thimm, S. Villate, F. Cerutti, N. Oren, H. Strass, and M. Vallati. The first international competition on computational models of argumentation. *AI Magazine*, 37(1):102–105, 2016.
- [Vega, 1990] W. Vega. Kernels in random graphs. *Discrete Mathematics*, 82:213–217, 1990.
- [Walicki and Dyrkolbotn, 2012] M. Walicki and S. Dyrkolbotn. Finding kernels or solving SAT. *Journal of Discrete Algorithms*, 10:146–164, 2012.
- [Wen *et al.*, 2016] L. Wen, K. Wang, Y. Shen, and F. Lin. A model for phase transition of random answer-set programs. *ACM Transactions on Computational Logic*, 17(3):Article 22, 2016.