Revista Iberoamericana de Argumentación

ἐπει δὲ ταύτην ἐπιστήμην ζητοῦμεν

Segunda Época

RIA 14 (2017): 1-28

Directores: Luis Vega y Hubert Marraud **Secretaria:** Paula Olmos

ISSN 2172-8801 / doi 10.15366/ria / https://revistas.uam.es/ria

Cumulative Arguments in Artificial Intelligence and Informal Logic

Argumentos acumulativos en inteligencia artificial y lógica informal

DOUGLAS WALTON
Centre for Research in Reasoning, Argumentation and Rhetoric
University of Windsor
Windsor, Ontario N9B 3P4 Canada
waltoncrrar@gmail.com

THOMAS F. GORDON
Institute of Computer Science
University of Postdam
August-Bebel-Str.89. 14482 Potsdam, Germany
thomas.gordon@fokus.fraunhofer.de

Artículo recibido: 03-02-2017
Artículo aceptado: 01-05-2017

ABSTRACT

We define a new type of argument structure specifically for modeling cumulative arguments and then show how this structure is general enough to simulate linked and convergent arguments. Argumentation schemes are associated with argument weighing functions in this language, where the weight of an argument can depend on the status (labeling) of its premises in an argument graph. Several key examples are used to illustrate the modeling of cumulative arguments, as well as linked and convergent arguments, with this approach. One hypothesis suggested by the analysis of these examples is that cumulative arguments can be treated in the same way as what is called argument accrual in artificial intelligence.

KEYWORDS: argument accrual, argument structure, argument weighing, artificial intelligence, convergent argument, cumulative argument, informal logic, linked argument.

RESUMEN

En este artículo se define un nuevo tipo de estructura argumentativa que caracteriza específicamente los argumentos acumulativos y se muestra que dicha estructura es, en general, suficiente para simular las estructuras tanto de argumentos dependientes como convergentes. Los esquemas argumentativos se asocian, por medio de este lenguaje, a las funciones de ponderación de argumentos, donde la fuerza de un argumento puede depender del estatus de sus premisas expresado por medio de una etiqueta en el diagrama argumentativo. Se ofrecen varios ejemplos significativos para ilustrar el diagramado de los argumentos acumulativos, tanto dependientes como convergentes, que propone este enfoque. Una de las hipótesis que surge de la consideración de dichos ejemplos es que los argumentos acumulativos pueden analizarse mediante lo que en inteligencia artificial se conoce como agregación de argumentos (*argument accrual*).

PALABRAS CLAVE: agregación de argumentos, argumentos acumulativos, argumentos convergentes, argumentos dependientes, estructura argumentativa, inteligencia artificial, lógica informal, ponderación de argumentos.

Servicio de Publicaciones de la Universidad Autónoma de Madrid

1. INTRODUCTION

There is broad agreement in the logic textbooks, and the literature in argumentation and informal logic generally, about how to apply argument structures to the natural language examples of argumentation in discourse. Such structures include the well-known single arguments, linked arguments, convergent arguments, and so forth. But in particular, another type of argument, that of cumulative arguments, evidence-accumulating arguments as they might be called, is comparatively less widely accepted and used, even though its importance has been acknowledged (Walton, 1996; Snoeck Henkemans, 2003). At the same time, a comparable notion, called accrual of arguments, is a growing subject in the literature on formal models of argumentation in artificial intelligence. One problem is to try to figure out whether what is referred to as cumulative arguments in informal logic and argumentation theory is the same as, or somehow closely related to, what is called accrual of arguments in artificial intelligence (AI). This paper presents a new computational model of cumulative argument and uses this model to shed some light on the relationship between cumulative arguments and argument accrual, to help the multidisciplinary field of argumentation studies move forward.

There are many examples of cumulative arguments in the literature that could be used, but obviously there is no space in a paper of this sort to try to analyze all of them in depth. Especially complex examples might be very interesting to analyze, but would also involve many other problems of analysis of natural language discourse that would be distracting. So in this paper we have selected some relatively simple examples of cumulative arguments with the purpose of trying to get some basic idea of how best to formally model them. Because the formal models of argument from artificial intelligence we present are clearly and precisely defined, they are useful to try to gain some clarity on moving forward with this class of arguments. One such example is prominently featured in this paper. Walton, Tindale and Gordon (2014) showed that the Carneades Argumentation System (CAS) can be used to model some ancient examples of cumulative argumentation, but needed to be extended to model important features of the ancient snake and rope example of Carneades the philosopher (214/3-129/8 BC), as described in the writings of Sextus Empiricus (Sextus Empiricus, 1938). The new version of CAS described in the present paper, CAS2, was developed explicitly to provide the features required to more fully handle cumulative arguments of the kind illustrated by the snake and rope example. In particular, the new version of CAS uses argument weighing functions, associated with argumentation schemes, to give arguments more or less weight, depending on the status of the premises of the argument. The argument weighing function of the argumentation scheme of an argument determines whether the argument is linked, convergent, cumulative or of some other kind.

2. ARGUMENTATION STRUCTURES

Four basic argument structures are widely recognized in the literature on argumentation in informal logic, linked arguments, convergent arguments, serial arguments and divergent arguments. In a linked argument, the premises must all be acceptable to support the conclusion. What is called a convergent argument can also be described as two separate arguments, each of which independently of the other supports the conclusion. A serial argument is one in which the conclusion of one argument is reused as a premise in the next argument, and so forth in a finite sequence forming what is often called a chain or sequence of argumentation. A divergent argument is one in which two different conclusions are derived from the same premise. Of these four, the first three are much more prominent in the literature on argumentation structures. The divergent structure tends not to be mentioned very much.

A brief survey in (Walton 1996, 133-134) explains how evidence-accumulating arguments have been treated in informal logic. Such arguments, often called cumulative arguments, have a characteristic that each premise alone lends some support to the conclusion, but with each additional premise the supported conclusion gets stronger and stronger. This pattern of argumentation, or structure of argumentation as it might be called, is extremely familiar to all of us. It could also be called updating of evidence for and against hypothesis as new evidence is collected in a given case. This form of evidential reasoning is highly familiar in both scientific and medical reasoning. For example, a doctor examining a patient might see some evidence, such as red spots on the patient's skin, a finding that might give a slight amount of evidence for a preliminary hypothesis that the patient has the measles. Further investigation might reveal supportive evidence, but on the other hand it might reveal additional evidence showing definitively that the patient's condition has some other cause than a measles infection. Strangely, although such a cumulative sequence of reasoning is clearly very common, practitioners of informal logic have been unable to reach any consensus on how to define it in a precise way.

The approach of The Amsterdam School to these argument structures appears to be quite different from the informal logic approach. In an extensive review of the scholarly literature and the textbooks in informal logic and argumentation, Freeman (2001) has shown that there is a pervasive confusion in argumentation studies on how to define and name the basic kinds of argument structures. He has shown that it is tempting to regard the distinction between linked and convergent arguments as coinciding with the distinction made by the pragma-dialectical school between multiple and coordinatively compound argumentation. However, he argues convincingly that this pair of distinctions should be regarded as marking something quite different. Indeed, Freeman (2001, 398) even showed that such differences go even deeper, because van Eemeren and Grootendorst used the term ‘argument’ in the same way that many North American scholars use the term ‘premise’. He also showed that this difference is disciplinary in nature because the former group approaches the concept of an argument from a dialectical point of view whereas the latter group tends to approach the concept of argument from a logical point of view. The result of this is that the current terminology on basic argument structures in the literature is deeply confusing.

This has created many problems, but one of special concern here is that there is no clear agreement on how to define the notion of a cumulative argument, even though this term has been widely used in the literature, and all indications are that scholars are using the term in sharply different ways. Freeman has provided highly extensive and useful survey of these differences, but has little to say specifically on differences concerning the meaning of the term ‘cumulative’. In fact he only mentions the term cumulative once (Freeman, 2001, 401). Hitchcock (2003) in a commentary on a paper of Snoeck Henkemans (2003) that builds on the valuable work of Freeman (2001), does discuss this term in relation to the usage of the pragma-dialectical school.

According to the definition proposed by Snoeck Henkemans (2003, p. 5) cumulative coordinative argumentation, as she calls it, consists of a number of reasons that each by themselves give some support to the standpoint and that should be sufficient when taken together to convince the antagonist of the acceptability of the standpoint. She adds that the force of the individual reasons may vary. In this definition, the term ‘reason’ is used instead of the term ‘argument’. Here too the terminology is unsettled (Prakken, 2005, p. 1). A reason could be the antecedent of a conditional forming one premise of an argument (the so-called warrant), or it could be a premise in an argument.

Hitchcock (2003, p. 2) has clarified this terminology by comparing the key terms in the two approaches. Figure 1 indicates how each term as used by the Amsterdam School can be equated with the comparable term in the informal logic lexicon.

doi: 10.15366/ria2017.14

Amsterdam (Pragma-Dialectics School)	Windsor (Informal Logic)
argument	premise
argumentation	argument
complementary coordinative argumentation	linked argument
cumulative coordinative argumentation	convergent argument
multiple argumentation	independent arguments
standpoint	conclusion
subordinative argumentation	serial argument

Table 1: Terminological Comparison of Argument Structure Names

Cumulative coordinative argumentation, equates to (or is perhaps a subspecies of) what is called the convergent type of argument structure in the informal logic approach. Hitchcock (2003) has shown that the way Snoeck Henkemans, and the pragma-dialectical school generally, use the expression 'cumulative coordinative argumentation', it is taken to be equivalent to what is normally called a convergent argument in informal logic. It will be clear from our paper that we use the term 'cumulative argument' in quite a different way and that it is important for us to start afresh. We accept the distinction between linked and convergent arguments, but we approach the notion of a cumulative argument in an open-minded way, leaving it to be interpreted in light of our own examples we use to illustrate its use in argumentation.

Let's take a look at an example (8), from Snoeck Henkemans (2003, 5), which is in the form of a small dialogue. At the first move, Paula tries to defend herself against Anton's criticism of her argument that a movie must be good because it is playing in the Cinecenter (a popular theater). He has argued that this does not guarantee that the movie will be good. She replies that Theo was very enthusiastic about it. The dialogue is quoted below from (Snoeck Henkemans, 2003, 5).

> Paula: It must be a good movie, because it is playing in Cinecenter.
> Anton: It's not as if I never saw a bad movie in Cinecenter.
> Paula: Yes, but Theo was also very enthusiastic about it.

Paula's reply is classified as an instance of cumulative coordinative argumentation. On Snoeck Henkemans' account, cumulative coordinative argumentation consists of a number of sub-arguments that each give some support to the claim at issue and that should be sufficient to prove the claim to the respondent when taken together.

The sequence of argumentation in this example starts out with an initial argument to support a claim, but then the other party expresses doubt about this argument. The response of the first party is to bring forward an additional argument that has the effect

doi: 10.15366/ria2017.14

of strengthening the initial argument. Even though the two arguments are not convincing enough by themselves, when taken together they may be sufficient to convince the other party. Such a secondary reason is described as an additional reason that is added to an initial argument. Presumably this procedure could go on until the dialogue is closed. A third reason could be given if the protagonist assumes that the two prior arguments are separately not convincing enough, and so forth. In such a case the term ‘cumulative’ is appropriate, because there is an accumulation of weight in support of the ultimate claim as each new argument is added to the sequence.

This example will fit with the general approach to cumulative argumentation presented in this paper, once it is established how the term ‘argument structure’ is to be defined. In what follows we will take an argument structure to be a subgraph of an argument graph, which could be visualized as an argument diagram. An argument diagram is essentially a graph structure. On this approach, an argument structure could be defined as a subgraph of a larger graph. So, for example, there might be a linked argument connected to a convergent argument in such a way that the two arguments together constitute a serial argument, but this serial argument is simply part of a larger network of argumentation leading to the ultimate claim at issue in a given case.

3. A TYPICAL EXAMPLE

A typical example of a cumulative argument is an argument that is brought forward, usually one that represents some piece of evidence that does not by itself weigh heavily as a strong argument, and then another piece of evidence is introduced that supports the evidential weight of the first argument. It is typical of such arguments that they can form a sequence. In such a case a network of argumentation containing single, linked and convergent arguments is produced in which there is a cumulative buildup for the ultimate conclusion.

A typical argument of this sort that is easy to grasp is the well-known example from the *Study in Scarlet* of Sherlock Holmes’ reasoning used to illustrate his famous method of using “deductive logic” to solve criminal cases (Walton, 1996). Watson had returned from Afghanistan, where he was wounded in a military campaign. He was interviewed by Holmes as a potential tenant to share the famous flat at 221B Baker Street. Holmes asked whether Watson had just been in Afghanistan, and Watson was surprised by such a lucky guess. Holmes replied that he knew that Watson came from Afghanistan and reconstructed the sequence of steps he used to arrive at this conclusion by logical reasoning (quoted from Walton, 1996, 99).

> Here is a gentleman of a medical type, but with the air of a military man. Clearly an army doctor, then. He has just come from the tropics, for his face is dark, and that is not the natural tint of his skin, for his wrists are fair. He has undergone hardship and sickness, as his haggard face says clearly. His left arm has been injured. He holds it in a stiff and unnatural manner. Where in the tropics could an English army doctor have seen much hardship and got his arm wounded? Clearly in Afghanistan.

This sequence of argumentation could be taken as the canonical example of cumulative argument, except that only pro arguments are considered. Yet it is typical of many common instances of scientific reasoning. Conan Doyle was a medical doctor, and quite familiar with the use of cumulative argumentation to build up a scientific chain of reasoning used to arrive at a diagnosis of a patient's illness, or to assemble a network of circumstantial reasoning in a powerful cumulative buildup of forensic evidence in a criminal case.

As each bit of evidence is brought in, based on Holmes' astute observations, the conclusion that Watson must have just been in Afghanistan is more and more strongly supported by the growing body of evidence that has been assembled. First, Holmes arrived at the conclusion that Watson must be an army doctor, based on two aspects of his appearance. Holmes then observes that Watson must have just come from the tropics, based on several observations about Watson's skin. The first argument supports the ultimate conclusion only slightly, but when the second argument is added to it, the conclusion is more strongly supported. Finally, Holmes introduces more evidence suggesting that Watson must be an army doctor, leading to his asking the question where an English army doctor could have seen so much hardship, and then he answers his own question by producing the conclusion that Watson must have recently been in Afghanistan.

In section 4 an argument diagram will be used to join the sub-arguments together and display them visually in a graph structure showing how the characteristic pattern of the argumentation is part of a cumulative buildup of evidence supporting the ultimate conclusion.

4. ACCRUAL OF ARGUMENTS

Verheij (1995, p. 217) addressed what he called an often overlooked problem of argument accrual by posing a question: "how do we deal with arguments that are on their own defeated but together remain undefeated?" To handle this problem, he introduced the notion of compound defeat of arguments, where groups of arguments can be defeated by other groups of arguments. This way of posing the problem suggests that the problem of argument accrual arises where you have a group of several arguments interacting with each other, for example in a typical argument diagram of a complex enough sort showing groups of arguments interacting with each other. Presumably, in such situations, an argument placed somewhere in the diagram can be proved or refuted several times as other arguments attacking it or supporting it are taken into account.

(Verheij, 1995, 217) explains his approach to accrual of arguments by offering the following simple example concerning three arguments, a1, a2 and a3. His definition concerns the accrual of arguments in defeasible argumentation, where a defeasible sequence of argumentation is defined as one in which a given argument can be supported or defeated as new arguments become available. In this sense, defeasible argumentation is open-ended, so to speak. The status of an argument as justified or not by the evidence can change as new evidence enters into consideration in a case at issue. The argument in the example has to meet three requirements: (1) the argument a1 defeats the argument a2, if a2 and a1 are the only arguments available, (2) the argument a1 defeats the argument a3, if a3 and a1 are the only arguments available, (3) but the arguments a2 and a3 join together to support the argument a1 where all the mentioned arguments are available.

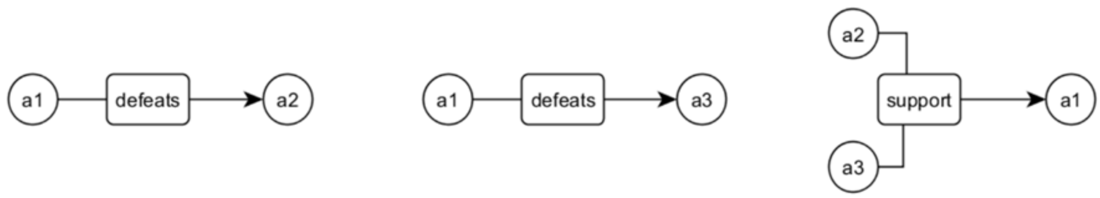

Fig. 1: Accrual of Arguments

In the left diagram, a1 defeats a 2. In the diagram in the middle, a1 defeats a3. In the right diagram, a2 and a3 support a1.

Prakken (2005, p. 2) offered a real example showing how arguments support or attack each other as part of a sequence of argumentation in which accrual takes place. In this example, two arguments are given not to go jogging, namely that it is hot and that it is raining but then he considers the possibility that a particular jogger might find the combination of heat and rain a pleasant combination. In this instance, the accrual is a weaker argument not to go jogging than the accruing reasons. He even suggests the possibility that the combination of heat and rain while jogging might be so pleasant to the jogger that it turns into a positive argument for going jogging. In this example, heat and rain each independently offer arguments against going jogging, but taken together each independent argument is weakened. Or it may even be the case when both are put together, the original conclusion is defeated.

Prakken (2005, p. 2) writes that the starting point of any attempt to formalize accrual is the principle that adding more arguments can make one's case stronger. He also formulates three general principles that help to define the notion of accrual he has in mind. The first principle is that an accrual can be weaker or stronger than the arguments being accrued, considered separately. In the jogger example, the combination of heat and rain may be pleasant when running, whereas heat or rain alone may be unpleasant. Hence in this case, the accrual (hot and raining) is a stronger argument pro going running than either hot or raining alone, each of which are weaker arguments con going running. The second principle is that when a larger accrual is applied, the application makes all of its lesser versions inapplicable. This is because each of the individual accruals for and against the claim is meant to consider all the available information at the point they were taken into account. However, once they are taken into account they can be cast aside because they only take part of the information into account. The third principle is that when an individual argument that was part of the accrual process turns out to be flawed, it does not take part in the accrual. For example, in a case of accrual of witness testimony evidence for a claim at issue, if one of the witnesses turns out to be incompetent, the argument from his testimony is undercut, and is cast aside in the accrual procedure.

This example and three principles suggest that as new evidence is brought into a complex network of argumentation, so that the argument diagram gets larger and larger, the original argument that was first considered may be either supported or attacked as more and more arguments connecting with it are taken into account. If this is what is meant by argument accrual, it appears to be very similar to, or perhaps even the same as the notion of cumulative argumentation that has appeared occasionally in the literature on argumentation and informal logic. But it is not easy to tell, because as Prakken shows, there are different approaches to the formal modeling of accrual in the AI literature.

5. DEVELOPMENT OF THE CARNEADES ARGUMENTATION SYSTEM

The original version of the Carneades Argumentation System (CAS1) modeled arguments as bipartite directed graphs, containing statement and argument nodes (Gordon, Prakken and Walton, 2007), represented in argument diagrams as rectangles and circles, respectively.

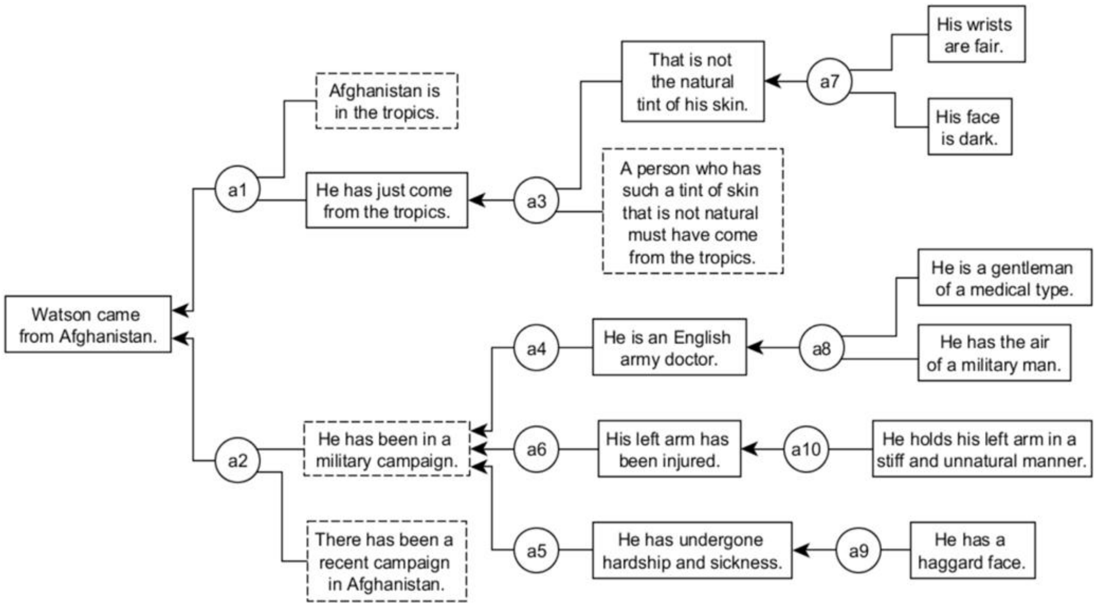

Figure 2: A CAS Argument Diagram of the Sherlock Holmes Example

In argument graphs, argument nodes are linked to their premises and conclusion nodes, which are statements (propositions). CAS1 had two types of arguments, pro and con, for arguments which support or attack their conclusion, respectively. The dashed boxes represent implicit premises. Figure 2 illustrates CAS1 with an argument graph representing an interpretation of the argumentation in the Sherlock Holmes example from section 2.

There are two main arguments in the example, a1 and a2, each providing a separate line of argumentation supporting the ultimate conclusion that Watson came from Afghanistan. Argument a1, by itself, is fairly weak, because there are other places, such as Africa, that are also in the tropics. However, argument a2 provides additional support by providing evidence that Watson had recently been in a military campaign. Taken together with the assumption that there had been a recent campaign in Afghanistan, a2 provides additional support for the conclusion that Watson came from Afghanistan. This example well illustrates how single, linked and convergent arguments can be combined to fit into a complex sequence of argumentation that also contains cumulative argumentation.

There have been four versions of the Carneades software. The first three versions were based on the formal model of argument from 2007 mentioned above (Gordon, Prakken and Walton, 2007), called CAS1 here. The first implementation was a command line tool written in 2006-2008. The second (2011) implementation is a desktop application, with a graphical user interface, sometimes called the Carneades Editor. Figure 2 is a typical argument diagram of the kind produced by version 2. The third version is a multi-user web-application, with a client-server architecture, developed in the European MARKOS project (2012-2015). The MARKOS project developed a prototype application for browsing and analyzing functional, structural and licensing properties of open source software. It includes a license analyzer tool, based on Carneades, that applies a formal representation of domain-dependent argumentation schemes for copyright law to facts in a repository about open source software projects, to automatically construct arguments about licensing issues and build an argument graph.

Version 4 is the current version of the Carneades software. It is based on a new formal model of argument (Gordon and Walton, 2016), called CAS2 here. CAS2 provides improved support for cumulative arguments, cyclic argument graphs, practical reasoning, and multi-criteria decision analysis. The source code of all four versions can be accessed on the Internet.[1] Carneades 4 is now online.[2] The present paper shows how CAS2, as implemented in version 4 of Carneades, provides a better way of evaluating cumulative arguments than CAS1, as implemented in earlier versions of the Carneades software, using the most famous ancient snake and rope example.

In the example shown in figure 2, there are only pro arguments supporting the conclusion. But it is also possible to have cases of con arguments appearing in the same chain of argumentation that detract from or go against the same conclusion. According to Hitchcock (2003, 3), the taxonomy of different types of argument structures needs to be supplemented by a category of cumulative arguments which recognizes an acknowledged kind of argument against the ultimate conclusion. He calls such arguments "counter-considerations". For example, scientific reasoning is supposed to be falsifiable, meaning that even if a conclusion is supported by the considerable buildup of positive evidence, it needs to be open to the possibility of being refuted by negative evidence. The jogging example, which can be analyzed using the tool presented in

[1] https://github.com/carneades

[2] http://carneades.fokus.fraunhofer.de/carneades

doi: 10.15366/ria2017.14

section 8 using CAS2, illustrates this feature.

6. THE SNAKE AND ROPE EXAMPLE

An ancient form of Platonism called academic skepticism claimed that certain (absolute) knowledge of the truth is impossible. More precisely, this modified form of skepticism claims that the standard of proof as knowing something beyond all doubt is unrealistic for human agents, or perhaps even for machine agents that have knowledge and use it to act autonomously. Carneades (c. 213 - c.128 B.C.), a leading academic skeptic, criticized the Stoic philosophers for claiming that some propositions can be known to be true beyond all possibility of doubt. Carneades was the head of the third Platonic Academy. He did not claim that we cannot have knowledge at all, but held what might be called a modified form of skepticism.

Carneades' theory of argumentation arose from skeptical doubts about the Stoic claim that a cognitive impression (a mental image) of an object that is clearly perceived provides an accurate grasp of the nature of the object so that the content of the impression can be accepted as a true proposition. Skeptics, however, attacked this claim using the familiar examples of deceptive appearances. But an objection to skepticism is that it lacks a criterion for acceptance in making rational decisions on how to act prudently in the daily affairs of life. To provide a basis to respond to this objection, Carneades put forward a theory of defeasible reasoning that can be used for this purpose. On this theory, initial impressions can be tentatively accepted, provided they are open to further evidential testing (based on other impressions) that can falsify them. On this view, we can have knowledge, but only a kind of defeasible knowledge that is continually open to testing based on new evidence.

The main sources of our knowledge about Carneades' theory are from the writings of Sextus Empiricus (c.160 - 210 CE), a physician and philosopher. The philosophical works of Sextus are the most complete surviving source of knowledge about Carneades' theory. Carneades' main example is that of the rope and snake, as reported by Sextus Empiricus (AL 188), quoted from the Loeb Library translation (Sextus Empiricus, 1938, 101-102).

> For example, on seeing a coil of rope in an unlighted room a man jumps over it, conceiving it for the moment to be a snake, but turning back afterwards he inquires into the truth, and on finding it motionless he is already inclined to think that it is not a snake, but as he reckons, all the same, that snakes too are motionless at times when numbed by winter's frost, he prods at the coiled mass with a stick, and then, after thus testing the presentation received, he assents to the fact that it is false to

suppose that the body presented to him is a snake.

The man sees what looks like a coil of rope in a dimly lit room, but because of the uncertainty of seeing it clearly in the dim room, he tentatively accepts the assumption that it could be a snake. Acting on a prudential concern for safety, he jumps over the object after entering the room. Looking back afterward, he sees that the object did not move. Given this new evidence, he adopts the new assumption that the object is a rope. Not satisfied however, he prods the object with a stick. Still it does not move. This additional test confirms the hypothesis that the object is a rope and not a snake.

On Carneades' theory of rational acceptance, as illustrated by the snake and rope example, an appearance can be accepted provisionally if it meets three criteria: (1) it appears to be true, (2) it can be accepted even more strongly if it is stable (Sextus, AL 176), meaning that is consistent with other propositions that appear to be true, and (3) it can be accepted even more strongly if it tested and passes the test. A corollary is that the proposition accepted by criterion 1 needs to be rejected if it fails to meet criterion 2 or 3. This theory is widely applicable. Sextus (AL 184) presented medical examples. He also gave the following example: "when we are investigating a small matter we question a single witness, but in a greater matter several, and when the matter investigated is still more important we cross-question each of the witnesses on the testimony of the others". This example suggests that Carneades' theory can be applied to legal applications of argument based on witness testimony evidence (Walton, 2008; Gordon, 2010).

7. USING CAS1 TO EVALUATE THE ARGUMENTS IN THE SNAKE AND ROPE EXAMPLE

Pollock (1995, 40) gave a famous example to draw a distinction between two kinds of refutations called rebutters and undercutters (Pollock, 1995, 40). A rebutter attacks the conclusion of a prior argument. An undercutter casts doubt on whether the claim holds by attacking the inferential link between the premises and conclusion of the prior argument. He used the red light example (1995, 41) to illustrate an undercutter. Suppose an object looks red to me, and I conclude it is red for this reason, but then I find out that it is illuminated by a red light. But I know that being illuminated by a red light can make an object look red when is not. This is not a reason for concluding that the object is not red. It might be red, after all. But it is a reason that undercuts the argument that the object is red simply because it looks red.

Some would say if I see a red object, its redness is immediately evident and would require no argument to support it. Some would say that in such a case, I am justified in claiming that I know the object is red, and that this claim is true beyond all doubt. However, what Pollock showed using the red light example is that my claim to see a red object in such a case is based on a species of defeasible reasoning of the following form: where φ is a perceptible property, an agent's having a φ image constitutes a prima facie reason for the agent to believe 'My circumstances exemplify φ'. This example teaches a very Carneadean lesson (referring to the Greek philosopher) about the defeasibility of empirical knowledge.

This form of reasoning from perception has been recognized as an argumentation scheme called argument from appearance. The form of this scheme has been expressed in various ways (Walton, 2006; Walton and Sartor, 2013), but here we choose a simple form of it AP (argument from perception) that is convenient for our purposes here.

(AP) If something looks like a type of object F, then it is an F.
This object O looks like a type of object F.
Therefore O is an F.

Applying this scheme to the example, if I see an object that looks red, I can reasonably draw the conclusion that the object is red. But if I find out that the object is illuminated by a red light the object might be red or might not. This new evidence undercuts the argument (to use Pollock's term) that it is red by undermining its support for the conclusion that the object is red.

Given this scheme, we can now begin to see how the snake and rope example can be modeled using CAS1. To show very simply how the CAS1 works, we will break the sequence of argumentation down into four steps. Let's begin with the first step.

The circular node in figure 3 contains the notation +AP, meaning that the given argument fits the scheme for argument for argument from appearance and is a pro argument. The proposition contained in the rectangle with the dashed border is marked as an implicit premise, an unstated proposition needed to make the given argument fit the scheme. The conclusion of the argument appears at the left, and the two premises supporting it are shown at the right. In CAS1, an argument is evaluated as justifying its conclusion if (1) the premises of the argument are accepted by the audience (2) the argument has not been undercut by any other arguments that defeat it and (3) the argument is strong enough to meet the standard of proof for its conclusion. Numerical weights can be attached to the argument representing the strength of the argument according to audience acceptance, represented as a fraction between zero and one. The user inputs this information. Statements (propositions) that have been accepted (assumed) or are acceptable (derived) are shown in boxes with a green background (light grey in the printed version).

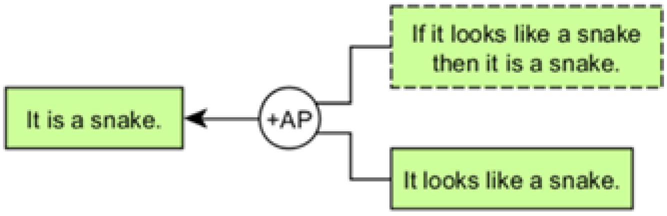

Figure 3: First Step in the Snake and Rope Example Evaluated Using CAS1

As shown in Figure 3, both premises of the argument have been accepted. CAS will then automatically derive the conclusion, which is thus also shown with a green background, provided that the three requirements above are met. This first step in the argumentation is very simple. It is a linked argument having the form of argument from appearance. There is no "audience" literally speaking, because the person in the example is not discussing what to do with another party. He is making this decision by himself, for all we know. But we are told that what he saw looked like a snake, and the other premise, the proposition in the box with the dashed outline, looks like it can be accepted as general common knowledge. Neither proposition is in dispute.

But the second step introduces a complication. It is an argument also based on the scheme for argument from appearance, but its conclusion is incompatible with that of the first argument. For the person in the example has to decide whether the object is a snake or a coil of rope. For purposes of deciding, he can't have it both ways. Otherwise, the second argument has the same basic structure as the first one, as shown in figure 4.

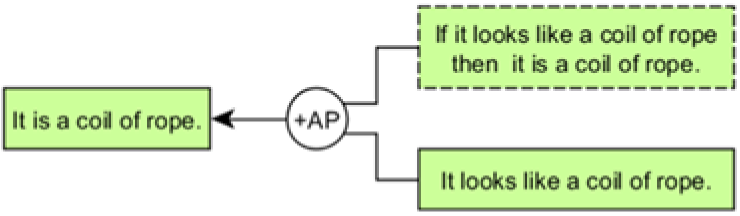

Figure 4: Second Step in the Snake Example Evaluated Using CAS1

At the next step, some new evidence enters the picture. At this state of the argumentation then, the person in the example being a rational agent, but one who apparently has some reason to enter the dark room, takes the precaution of jumping over the object as he passes into the room.

The object did not move when the person jumped over it. Hence we are taken to the next step in the sequence of argumentation. This argument step is shown in figure 5 as an instance of the scheme for argument from evidence to a hypothesis (EH).

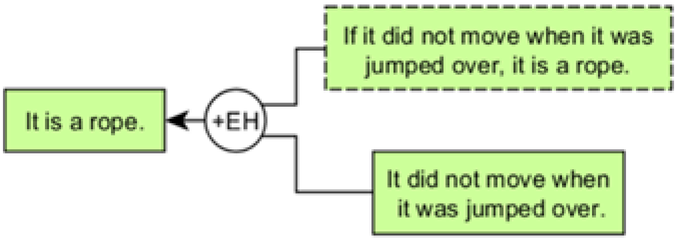

Figure 5: Third Step in the Snake and Rope Example Evaluated Using CAS1

Finally we proceed to the fourth and last argumentation step in the sequence. The person still wants to test his working hypothesis that the object is not a snake, so he prods the object with a stick to see if it will move. As shown in figure 6, it did not move. Hence CAS1 colors the conclusion box green.

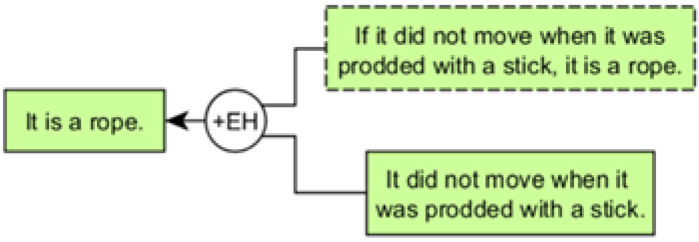

Figure 6: Fourth Step in the Snake and Rope Example Evaluated Using CAS1

CAS1 automatically draws the conclusion once again that the object is a rope. The final outcome then is the conclusion that the object was a rope, or at least this is the conclusion drawn based on the total body of evidence given.

Based on the text of (Sextus Empiricus, 1938, AL 184) just after the part of the snake and rope example quoted, it could be argued that, in the example, the first hypothesis that is entertained is that it is a snake because it is safer to assume it, not necessarily because it is the most plausible hypothesis. Walton, Tindale and Gordon (2014) use a scheme for argument from danger to model this aspect of Carneades' argument. So this contextual factor can be taken into account, even though it has not been considered here.

What we have shown is that CAS1 can be used to evaluate a sequence of cumulative argumentation, but in order to do so the user has to manually break the task down into a sequence of sub-evaluations containing several steps. As the sequence of argumentation moves from one step to the next, an evaluation of the next part of it can be made, showing how, at that step of the procedure, the ultimate conclusion is proved by the premises or not. However in utilizing this kind of procedure, the premises change at each move, as new evidence enters the picture. What is not possible in CAS1 is to put all of the evidence and arguments together into a single graph and have it infer that the object is presumably a rope, rather than a snake. CAS2 overcomes this limitation. Let's now see how.

8. THE CAS2 FORMAL MODEL OF CUMULATIVE ARGUMENTS

This section provides an overview of the new CAS2 formal model of structured argument, presented in detail in (Gordon and Walton, 2016). We begin by letting L be a logical language for expressing propositions. In CAS2, an argumentation scheme is defined as a tuple (e, v, g), where e is a function for weighing arguments which instantiate a scheme, v is a function for validating arguments, to test whether they properly instantiate an argumentation scheme, and g is a function for generating arguments by instantiating the scheme. For the purposes of this paper, we will only be concerned with evaluating (weighing) the argument, but it is necessary to understand validation as well. The validation function tells us whether the argument properly instantiates some scheme, including whether any premises required to instantiate the scheme are implicit (enthymemes). Given a set of schemes, we can apply each of their validation functions to some argument, to find which schemes are instantiated, if any, by the argument. These validation functions place restrictions on the premises and conclusion which must be satisfied in order for some argument to be an instance of the scheme.

An argument is defined as a tuple (s, P, c, u), where s is the scheme instantiated by the argument; P, a finite subset of L, is the set of premises of the argument; c, a member of L, is the conclusion of the argument; and u, a member of L, is the undercutter of the argument. There are three ways you can attack an argument. You can attack one or more of the premises, you can attack the conclusion, or you can attack the inferential link joining the premises to the conclusion (by arguing that some exception applies, for example). The last mode of attack is called undercutting (Pollock, 1995).

Both CAS1 and CAS2 model standards of proof for use as part of the procedure for evaluating arguments (Gordon, Prakken and Walton, 2007; Gordon and Walton, 2016).

- Scintilla of Evidence
  - There is at least one applicable argument
- Preponderance of Evidence
  - The scintilla of evidence standard is satisfied, and
  - the maximum weight assigned to an applicable pro argument is greater than the maximum weight of an applicable con argument.
- Clear and Convincing Evidence
  - The preponderance of evidence standard is satisfied, and
  - the maximum weight of applicable pro arguments exceeds some threshold α, and
  - the difference between the maximum weight of the applicable pro arguments and the maximum weight of the applicable con arguments exceeds some threshold β.
- Beyond Reasonable Doubt
  - The clear and convincing evidence standard is satisfied, and
  - the maximum weight of the applicable con arguments is less than some threshold γ.

Note that with this way of defining the standards of proof, the thresholds α and γ are left open to the user to put in, and are not given fixed numerical values by the formal model. The default standard which is automatically set is that of the preponderance of the evidence, but the user can change to one of the other standards if he or she wishes.

An issue is defined as a tuple (O, f), where O is a finite subset of L, representing the options (also called positions) of the issue. f is the proof standard of the issue, a function which tests whether an option satisfies its standard. A distinctive feature of CAS2 is that the issues need not be Boolean. There can be zero or more alternative options (positions) for each issue. This is useful for overcoming false dilemmas ("Have you stopped beating your spouse?") as well as for supporting deliberation dialogues, where the pros and cons of any number of options are compared.

Argument graphs in CAS2 are now tripartite, rather than bipartite, with nodes for statements, arguments and issues. More formally, in CAS2 an argument graph is defined as a tuple (S, A, I, R), where S, the statements of the argument graph, is a finite subset of the language L; A, the assumptions, is a subset of S assumed to be provable; I is a finite set issues, where every position of every issue is a member of S and no s ϵ S is a position of more than one i ϵ I,; and R is a finite set of arguments, in which all conclusions, premises and undercutters are members of S.

Argument diagrams in version 4 of Carneades, designed to support CAS2, have been extended with a new node type, diamonds, for representing issues. There can be any number of issues in a single diagram. Rectangles and circles are used, as before,

doi: 10.15366/ria2017.14

to represent statements and arguments, respectively.

Argument evaluation is carried out in a broadly comparable way in both CAS1 and CAS2. The evaluation process labels the statements in, out or undecided. In argument diagrams, in, out and undecided statement rectangles are shown filled with green, red and white backgrounds, respectively. Intuitively, a statement is in if and only if it has been assumed to be acceptable to a rational audience, or has been derived from such assumptions via the application of the arguments, argument weighing functions and proof standards. Roughly speaking, to say that a statement is out means that it is neither assumed nor supported by arguments and thus should be rejected by a rationale audience. A statement is undecided if it is neither in nor out. A statement can be undecided if there are cycles in the argument graph which cannot be resolved.

Argument weighing functions can be used to model linked, convergent and cumulative arguments as follows: A linked argument has weight 1.0 if all its premises are in, but otherwise has weight 0.0. A convergent argument has weight 1.0 if some premise is in, but otherwise has weight 0.0. The weight of a cumulative argument is the percentage of its premises that are in.

Two differences between CAS1 and CAS2 to be aware of concern their handling of issues and con arguments. In CAS1 all issues are Boolean and implicit. Every statement node implicitly represents a Boolean issue. If a statement is in, its complement (negation) is out and vice versa. In CAS2, on the other hand, issues are explicit, not limited to two options, and are represented with a new node type (diamonds) in argument diagrams. In CAS1 con arguments were modeled explicitly, as a type of argument. In CAS2, on the other hand, con arguments are modeled implicitly, as arguments pro another position of the same issue. Put differently, in CAS2 there are only pro arguments for different positions of issues. Since at most one position of an issue may be in, an argument pro one position is also, implicitly, an argument con every other position of the same issue.

9. USING CAS2 TO EVALUATE THE ARGUMENTS IN THE SNAKE AND ROPE EXAMPLE

How the snake and rope example is evaluated as an argument can now be shown in the argument diagram in figure 7. In Carneades 4, which implements the CAS2 model, statements which have been assumed to be acceptable (i.e. assumptions) are shown in argument diagrams with the text underlined, and statements which are in are shown with

doi: 10.15366/ria2017.14

a green background, as before. Statements which are out are shown with a red background (darker grey in a printed version). Assumptions are always in, so in addition to their underlined text, they are also shown with a green background. This improvement makes it possible to distinguish whether a statement is in (green) because it has been assumed (underlined) to be acceptable or because it has been derived via arguments to be acceptable (not underlined). The assumptions are determined the user on behalf of the audience. They represent the statements the user expects would be accepted by the audience without argument.

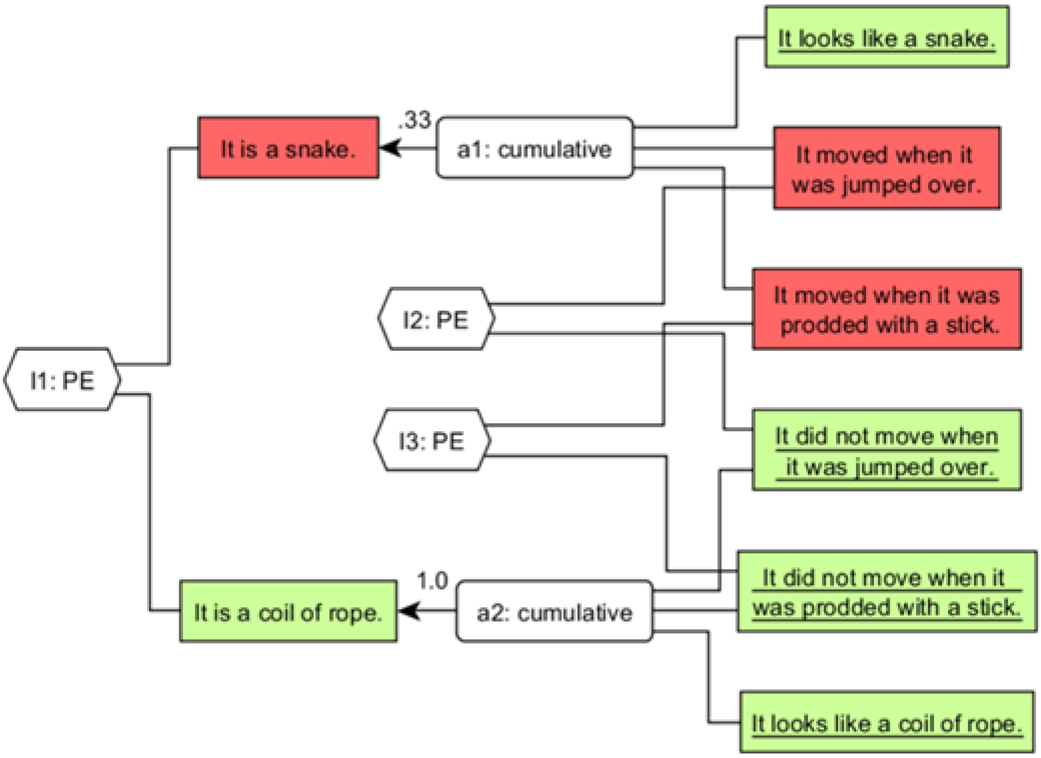

Figure 7: The Snake and Rope Example Evaluated Using CAS2

In figure 7, the ultimate issue, labeled as I1, is whether the object is a snake or a coil of rope. But there are also two sub-issues. Sub-issue I2 is whether the object moved when it was jumped over. Sub-issue I3 is whether it moved when it was prodded with a stick. In the first cumulative argument, a1, shown at the top right of figure 7, the statement for the premise that it looks like a snake has been assumed and is therefore in, whereas the other two premises are out. Thus a1 has a weight of 0.33 (1/3). However in the bottom cumulative argument, a2, all three premises are in. Therefore a2 has a weight of 1.0 (3/3). Hence, when applying the preponderance of evidence standard of issue I1, the option that it is a coil of rope satisfies the standard, because it is a supported by an argument, a2, which weighs more than any argument supporting the other option of the issue, that it is a snake. Hence the "It is a coil or rope." statement is in and shown with a green background, while the "It is a snake." statement is out and shown with a red background.

We suggest that at this point the reader might look back to figures 3 and 4 for a contrast with how the snake and rope example was modeled using CAS1. In figure 3, the conclusion was that the object is a snake. In figure 4, the conclusion was the object is a coil of rope. What is not shown in this way of modeling the arguments, using CAS1, is the conflict between these two propositions. But this is now clearly shown in figure 7 using CAS2. The two propositions are now modeled as two options for resolving a single issue.

10. THE JOGGING AND THE SHERLOCK HOLMES EXAMPLES USING CAS2

In this section, we explain how the jogging and Sherlock Holmes examples can be modeled using CAS2. Figure 8 shows a CAS2 argument diagram for the jogging example.

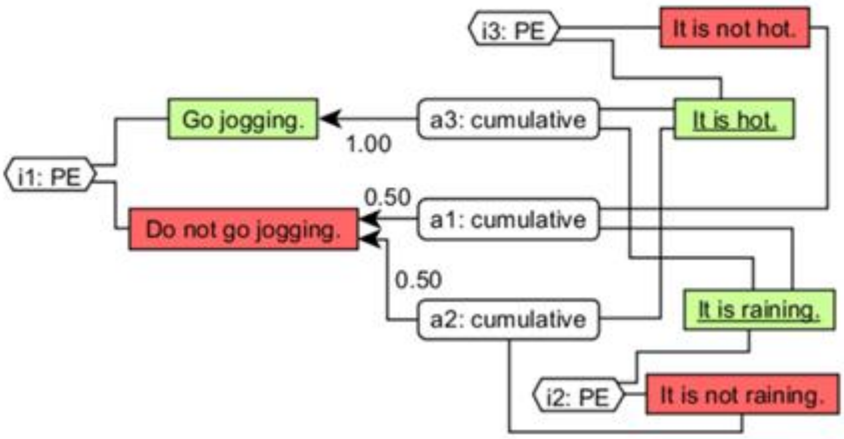

Figure 8: The Jogging Example Evaluated Using CAS2

Two of the premises on the right, for the statements that it is raining and it is hot, are shown in green rectangles and underlined, indicating that it has been assumed the audience accepts both of them. Since both of the premises of the cumulative argument a3 are in, a3 has the weight of 1.0 (2/2). For each of the cumulative arguments a1 and a2, on the other hand, only half of their premises are in. Thus a1 and a2 each weigh only 0.5 (1/2). Thus, only the "Go jogging." option of issue i1 satisfies the preponderance of evidence standard, since this is the only option which is supported by an argument, a3, which weighs more than any argument supporting the option to not go jogging.

Next we show how the Sherlock Holmes example can be modeled using CAS2. This is shown in figure 9. For those familiar with argument diagrams and argumentation and informal logic, this way of modeling the argumentation looks more familiar, and perhaps more easily suggests the kind of argumentation generally called cumulative in those fields.

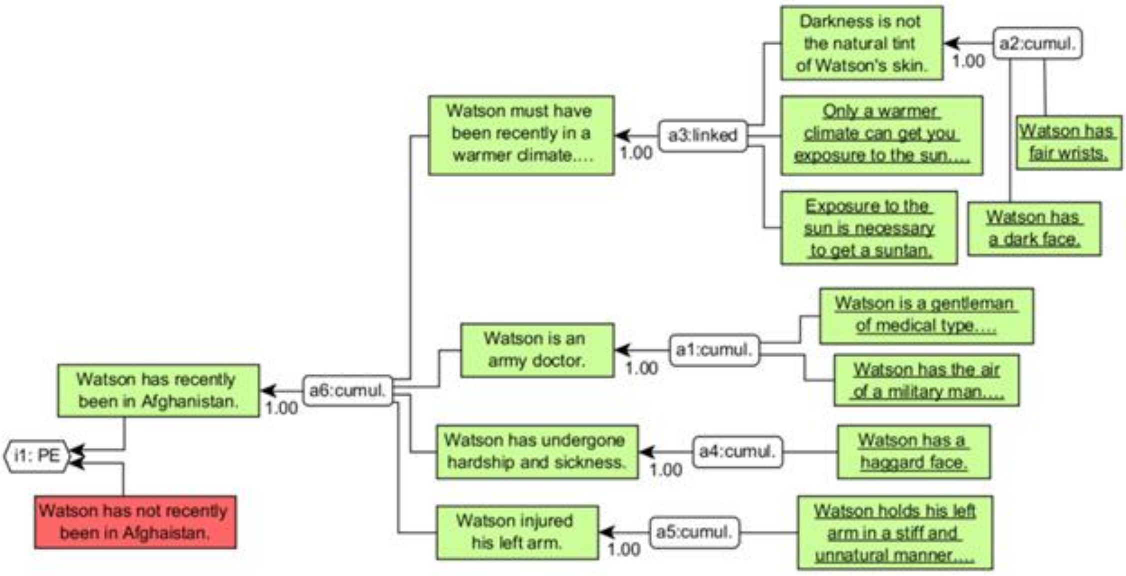

Figure 9: The Sherlock Holmes Example Evaluated Using CAS2

Note that in the CAS2 version of the Sherlock Holmes example, the issue to be resolved by the preponderance of evidence standard is whether Watson recently came from Afghanistan or not, and six arguments are used to support the claim that he recently came from Afghanistan. All of the arguments are cumulative, except a3. Two of the cumulative arguments, a4 and a5, only have one premise, each, so evidence is not yet accumulated. We have modeled them as cumulative arguments nonetheless, because they both present a collection of evidence to support a conclusion, even if this collection currently consists only of one piece of evidence each.

11. CONCLUSIONS

In the field of AI there are different approaches to modeling the kind of argumentation illustrated by these examples. Prakken (2005, 85) distinguished between two approaches to modeling argument accrual. The first, called the knowledge representation approach, models the accrual of arguments using a conditional with a conjunction of the propositions representing the new items of incoming evidence. He cites as examples of this approach the work of Pollock (1995) and the literature on probabilistic networks. The second, which he calls the inference approach, regards accrual as a step in a broader inference process. In this process all the relevant arguments are collected, and represented in some structures such as an argument diagram or argument graph, and then the whole conglomeration of arguments in the graph is aggregated by a weighing mechanism that decides the conflict when there are two conflicting sets of arguments. Prakken cites the model of cumulative argumentation of Verheij (2005) as an example of this approach. The argumentation model of CAS2 presented here is another example of the inference approach. It uses proof standards and argument weighing functions to aggregate conflicting arguments and choose among alternative options of issues.

CAS2 uses argument weighing functions, associated with argumentation schemes, to reinterpret linked, convergent and cumulative arguments as classes of argumentation schemes. We are not proposing that all validation functions for argumentation schemes accept any premises whatsoever. Going this far would eliminate the need for validation functions, since every argument would always be valid. Rather what we are proposing is that there are classes of argument which have some weight even if some of their premises are not acceptable, and that these kinds of arguments can be handled by a generalization of the concept of an argumentation scheme which uses weighing functions to weigh arguments instantiating these schemes. This generalized concept of an argumentation scheme can model linked, convergent and cumulative arguments as special cases.

The snake and rope example was used to illustrate how the new argument weighing feature of CAS2 can be used to model cumulative arguments. In CAS1, arguments were weighed manually by the audience, with no restrictions. Now, with CAS2, we are proposing to make argument weighing a further function of schemes, in addition to validation and generation. In CAS2 the concept of a scheme is generalized to include weighing functions, without or without further restrictions on the premises. In this new system, linked, convergent and cumulative arguments can all be reinterpreted as a special kind (class) of argumentation scheme.

We have examined the relationship between argument accrual and cumulative arguments, by showing how our formal model of cumulative arguments can also be used to reconstruct examples of argument accrual, in particular the jogging example. On this basis, we propose the hypothesis, subject to further research, that the expression 'cumulative argumentation' that is found in the literature on argumentation studies can be taken as essentially equivalent to what is meant by the expression 'accrual of arguments' in the AI literature. Inclusively speaking, all are instances of defeasible argumentation, where an accumulation of evidence needs to be taken into account as new evidence enters into consideration that is relevant to the line of argumentation as it moves towards an ultimate claim to be proved. Defeasible arguments are defined by Verheij (1995) and Prakken (2005) as arguments whose conclusion can be supported or refuted by further arguments put forward as the dialogue progresses.

We also want to point out a distinguishing characteristic of our approach, which can be illustrated by the following example: It must be a good movie, because: (a) It is playing at the Cinecenter, (b) it is made by a famous director, and (c) It has been compared to the Oscar-winning hit Moonlight. Analyzed separately, (a) is argument from sign. (b) is a causal argument scheme (from creator to the object created). (c) is an argument from analogy. In prior approaches to cumulative arguments in pragma-dialectics and informal logic, these three separate arguments would be joined together into a cumulative argument, but in such a way as to preserve the separate arguments. In this paper we have taken a different approach. In our formal model of cumulative argument, the various reasons for some conclusion are joined together as premises of a single cumulative argument, by instantiating a domain-dependent argumentation scheme for making decisions of the kind at issue, such as deciding whether or not a movie is worth watching, as in this example. That is, for each issue, a scheme is designed which specifies the relevant factors and dimensions of the problem and how they are to be weighed and aggregated when evaluating options, in the style of multi-criteria decision analysis.

Here we respond to two objections. The first one is that the examples of cumulative arguments we have presented seem to be examples of induction. If so, what is the difference between cumulative arguments and inductive reasoning? Here is our reply. Induction is not about drawing specific conclusions from the accumulation of different kinds of evidence, but rather about drawing general conclusions from the accumulation of the same kind of specific evidence: it rained the day before yesterday and it rained yesterday, therefore it will always rain. The first three balls in the box were white, so all the balls are white. Compare these with the snake and rope example, where the first piece of evidence was it didn't move when prodded, the second was it didn't move when jumped over and the conclusion was that the object is a rope. The premises

doi: 10.15366/ria2017.14

are different kinds of evidence and the conclusion is not a generalization. The snake and rope example could be recast as an inductive argument, but with different premises and a different conclusion: it didn't move when it was prodded twice, so it will not move no matter how many times it is prodded. Such an argument doesn't lead to the conclusion that it is a rope.

The second objection claims that the Sherlock Holmes example probably illustrates abductive reasoning, since Conan Doyle, as a medical doctor, was familiar with medical diagnostic reasoning of a kind often associated with abductive inference (Josephson and Josephson, 1994). In response it needs to be noted that Walton (2016, chapter 1) analyzed the argumentation in the Sherlock Holmes example both using the CAS1 approach, without cumulative arguments, as well has also in an alternative analysis that represents this same sequence of argumentation as a series of instances of abductive reasoning. This is interesting, but beyond noting it, we do not want to get drawn into a discussion of abduction as it would be too extensive and carry the paper away into a controversial area that is not directly germane to the main concern of the paper, about the meaning of cumulative arguments and how to formally model this meaning. Let it be said that there can be many ways of analyzing the argumentation in the Sherlock Holmes example, and examples of medical diagnostic reasoning (Josephson and Josephson, 1994).

We concede that our model offers no way to aggregate several arguments, instantiating various schemes, into a cumulative argument. Rather, in our approach a cumulative argument instantiates a single scheme, just as other kinds of arguments instantiate schemes. Thus, in the movie-going example, our approach would define a single scheme for going to movies, with premises for each of the factors to consider (location, director, comparisons with other movies, etc.) and then apply this scheme to construct/invent and weigh cumulative arguments for choosing movies to view.

12. FUTURE RESEARCH AND APPLICATIONS

We leave it up to others in future research to try to find a better model of cumulative arguments which can join together separate arguments without sacrificing the advantages of our approach. In particular, our approach has the advantage of generalizing the normative function of schemes, as acceptable patterns of reasoning, to cover domain specific, multi-criteria decisions, assuring the decisions of the same kind are made systematically, applying the same evaluation criteria and relative weights.

Future research needs to take up this challenge, but as this paper has shown, it seems to be such a wide-ranging and difficult challenge that it will need considerable work to make progress on solving it. It is to be hoped that the examples treated in our paper will help to move us some way forward to reaching a better consensus on this important methodological issue.

The ultimate test that will determine the outcome of future research depends on the analysis of examples such as the ones studied in this paper. But the examples also illustrate interesting variations. The jogging example is an instance of deliberation dialogue in which a jogger is trying to decide whether to go out jogging on a particular day in a region where the weather is continually changing, and where new information about the weather will affect his decision one way or the other. The snake and rope example surely involves some elements of deliberation, because the person entering the dark room wants to take the value of safety into account when making his decision on whether to enter the room or not, based on his seeing an object in the room that could be dangerous - it might be a snake. The narrow question addressed in the argumentation of this example is epistemic, as in information-seeking dialogues or critical discussions, since the person needs to evaluate the incoming evidence as he goes along in the cumulative sequence of making observations, to evaluate pro and con arguments to determine whether the object he sees is a snake or not. However this epistemological question about whether the object is a snake or rope can be understood to be a sub-issue of the practical question about whether or not to enter the room. That is, the information seeking dialogue or critical discussion may be embedded in a deliberation dialogue. The appropriate method to use to weigh and aggregate the cumulative arguments may differ depending on the purpose of the argumentation and the type of dialogue. In deliberation dialogue, the weighing functions need to take risks and benefits into account. In epistemological dialogues, we are more interested in finding the truth than in minimizing risks. Despite these differences, both kinds of dialogues share the problem of updating hypotheses using cumulative argumentation as further evidence becomes available.

The Carneades project currently underway is analyzing a number of other key examples along lines comparable to the way the snake and rope example of cumulative argumentation has been analyzed and evaluated above. Two areas of application are especially worth commenting on.

One of the most visible properties of scientific argumentation is that it proceeds by the method of putting forth a tentative hypothesis based on argument from the explanation of a given set of facts, and then proceeds to test the hypothesis by bringing in new evidence that can either support or refute it. This procedure clearly represents a defeasible sequence of argumentation based on two argumentation schemes, the scheme for argument to the best explanation and the scheme for argument from evidence to a hypothesis. What stands out as the central characteristic of any example of such a sequence is that the argumentation is cumulative. The evidence that comes in can have one of two effects - it can support the hypothesis or it can tend to show that the hypothesis is untenable. These are basically the two kinds of cumulative arguments. The snake and rope example in this paper illustrates how the CAS2 model of structured argument is able to handle both kinds of cumulative argument.

One of the most visible properties of legal argumentation, displayed clearly in civil or criminal trials, is the buildup of the mass of evidence as each side continues to bring in arguments relevant to the ultimate issue. This procedure too clearly exhibits cumulative argumentation, and the ultimate issue is decided by weighing the arguments on both sides as the new evidence comes in to be considered and evaluated. There are many differences concerning the protocols governing the dialectical structure of scientific argumentation versus legal argumentation, but what is basically evident is the central importance of cumulative argumentation in both kinds of cases. For these reasons, we suggest that the availability of a formal and computational model of argument that can handle cumulative argumentation is potentially very widely applicable in many significant domains.

REFERENCES

Freeman, J. B. (2001). Argument Structure and Disciplinary Perspective, *Argumentation*, 15(4), 397-423.

Gordon, T.F. (2010) An Overview of the Carneades Argumentation Support System. In C.W. Tindale and C. Reed, *Dialectics, Dialogue and Argumentation. An Examination of- Douglas Walton's Theories of Reasoning*, 145-156. College Publications.

Gordon, T. F., H. Prakken, and D. Walton, D. (2007). The Carneades Model of Argument and Burden of Proof. *Artificial Intelligence*, 171, 875–96.

Gordon, T. F. and Walton, D. (2016). Formalizing Balancing Arguments*. Proceedings of the 2016 Conference on Computational Models of Argument (COMMA 2016)*, 327–338. IOS Press.

Hitchcock, D. (2003). Commentary on Snoeck Henkemans, *OSSA Conference Archive*: https://www.google.ca/webhp?sourceid=chrome-instant&ion=1&espv=2&ie=UTF-8#q=david+hitchcock+Commentary+on+Snoeck+Henkemans

Josephson, J. R., & Josephson, S. G. (1994). *Abductive inference: Computation, philosophy, technology*. Cambridge: Cambridge University Press.

Pollock, J. L. (1995). *Cognitive Carpentry*. Cambridge, Mass.: The MIT Press.

Prakken, H. (2005). A Study of Accrual of Arguments, with Applications to Evidential Reasoning. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law*. New York: ACM Press, 85–94.

Sextus Empiricus (1938). *Against the Logicians*. Cambridge: Loeb Classical Library. Harvard University Press.

Snoeck Henkemans A. F. (2003). Indicators of Complex Argumentation. *OSSA Conference Archive*: http://scholar.uwindsor.ca/cgi/viewcontent.cgi?article=1566&context=ossaarchive

Van Eemeren, F., Houtlosser, P. and Snoeck Henkemans, A. F. (2007). *Argumentative Indicators in Discourse*. Berlin: Springer.

Verheij, B. (1995). Accrual of Arguments in Defeasible Argumentation. Dutch/German Workshop on Nonmonotonic Reasoning. *Proceedings of the Second Workshop*, pp. 217-224. Delft University of Technology, Utrecht University. Also published as report SKBS/B3.A/95-01. http://www.ai.rug.nl/~verheij/publications/dgws95.htm

Walton, D. (1996). *Argument Structure: A Pragmatic Theory*. Toronto: University of Toronto Press, 1996.

Walton, D. (2006). Argument from Appearance: A New Argumentation Scheme, *Logique et Analyse*, 195, 319-340.

Walton, D. (2008). *Witness Testimony Evidence: Argumentation, Artificial Intelligence and Law*, Cambridge, Cambridge University Press.

Walton, D. (2016). *Argument Evaluation and Evidence*. Cham: Springer.

Walton, D. and Sartor, G. (2013). Justification of Argumentation Schemes, *Argumentation*, 27(2), 2013, 111-142.

Walton, D., Tindale, C, W. and Gordon, T. F. (2014). Applying Recent Argumentation Methods to Some Ancient Examples of Plausible Reasoning, *Argumentation*, 28(1), 85-119.

D. WALTON: Douglas Walton (Ph.D. University of Toronto, 1972) is Distinguished Research Fellow of CRRAR (Centre for Research in Reasoning, Argumentation and Rhetoric) at the University of Windsor. He has been Visiting Professor at Northwestern University, University of Arizona, and University of Lugano (Switzerland). He is co-editor of the Critical Argumentation textbook series for Cambridge University Press. In 2011 he was Fernand Braudel Research Fellow of the European University Institute in Florence, where he collaborated on research in legal argumentation with Prof. Giovanni Sartor of the EUI and the Faculty of Law at the University of Bologna. In 2010 he was appointed to the Editorial Board of the journal Artificial Intelligence and Law. In 2009 he was given the Faculty of Arts and Social Sciences Dean's Special Recognition Award of the University of Windsor, in recognition of excellence in research, scholarship and creative activity.

T.F. GORDON: Tom Gordon conducts research on argumentation technology in the fields of "artificial intelligence and law" and "computational models of argument" at Fraunhofer FOKUS in Berlin, Germany, and holds an honorary professorship for argumentation technology at the Institute of Computer Science of the University of Potsdam. The goal of this research is to provide software tools for governance tasks in the policy life cycle: agenda setting, analysis, policy development and legislation, implementation (including administrative rulemaking and electronic service delivery), and monitoring.