ARGUMENT MINING BY APPLYING ARGUMENTATION SCHEMES

This paper builds a method to help a student of informal logic to go through a text of discourse in a natural language and identify common types of arguments that occur in the text. It is shown how this procedure is very helpful for students learning informal logic skills, as they sometimes misidentify arguments. The paper presents the state-of-the-art on what resources are available to build a useful argument identification procedure, and includes a survey of work done on automated argument mining tools used in artificial intelligence.

The aims of this paper are to survey the resources available for the project of building an exact method that will be helpful for the purpose of identifying arguments in natural language discourse, and to formulate some specific problems that need to be overcome along the way to building the method. It is argued that such a method would be useful as a tool to help students of informal logic identify arguments of the kind they encounter in natural language texts, for example in newspapers, magazines or on the Internet. The method proposed is based on the use of argumentation schemes representing common types of defeasible arguments [15] [19]. The idea is that each scheme is associated with a set of identifiers, (key words and markers locating premises and conclusions), and when the right grouping of identifiers is located at some place in a text, the argument mining method locates it as an instance of an argument of some particular, identifiable type (from a list of schemes).

The project is related to the development of argumentation systems in artificial intelligence. One of these technical initiatives, outlined in section 7, is the project of building an automated argumentation tool for argument mining. The idea is that this tool could go onto the Internet and collect arguments of specifically designated types, like argument from expert opinion for example. These technical initiatives are connected to the aim of finding an exact method for argument identification in informal logic, because the most powerful method would likely turn out to be to combine both tasks. The most powerful method would have human users apply the automated tool to identify arguments on a tentative basis in a text, and then correct the errors made by the automated tool. It is not hard to see how even a semi-automated procedure of this kind would be extremely helpful for teaching courses in informal logic.

As teachers of logic courses well know, judging whether an argument in a given text of discourse fits some abstract form of reasoning is a sophisticated task that many beginning students in courses in argumentation and informal logic have recurring problems with. Such courses are based on the identification, analysis and evaluation of examples of arguments found in magazines, newspapers, and the Internet, or whatever other sources of text materials are available. It is necessary, in order to do an adequate job of teaching an informal logic course, to have access to examples of commonly used arguments, and especially types of arguments that tend to be associated with common fallacies, like arguments from expert opinion, ad hominem arguments, appeals to force and threats, and so forth. To improve informal logic methods, we need to have systematic criteria for identifying a particular type of argument, like an argument from expert opinion, as found in a natural language text of discourse. For as we all know, natural language discourse is full of vagueness and ambiguity, and it can be very hard to pin down a real instance of some text to see whether it fits any abstract structure like a form of argument. Having a procedure for assisting with this task is simply a continuation of the kind of work that is being done every day in teaching courses and writing textbooks

in the field of informal logic. However, more exact methods would enable us to find new examples more easily, and to document and store them so they could be easily re-used

1. Teaching Students of Informal Logic to Identify Arguments

At the beginning, there are two specific tasks that need to be separated. One is the task of identifying arguments as entities that are distinct from other kinds of entities that occur in natural language discourse, like explanations. This is the task of distinguishing between arguments and non-arguments. This task is far from trivial, as verbal indicators are often insufficient to distinguish between something that is supposed to be an argument in something that is supposed to be an explanation [13]. The other task is that of identifying specific types of arguments. The earlier book on argumentation schemes [15] identified and described twenty-nine commonly used schemes that represent types of arguments familiar to anyone with a beginner's knowledge of informal logic.

1. Argument from Analogy
2. Argument from a Verbal Classification
3. Argument from Rule
5. Argument from Exception to a Rule
6. Argument from Precedent
7. Practical Reasoning
8. Lack of Knowledge Arguments
9. Arguments from Consequences
10. Fear and Danger Appeals
11. Arguments from Alternatives and Opposites
12. Pleas for Help and Excuses
13. Composition and Division Arguments
14. Slippery Slope Arguments
15. Arguments from General Acceptance
16. Argument from Commitment
17. Arguments from Inconsistency
18. Ethotic Ad Hominem
19. Circumstantial Ad Hominem
20. Argument from Bias
21. Ad Hominem Strategies to Rebut a Personal Attack
22. Argument from Cause to Effect
23. Argument from Effect to Cause
24. Argument from Correlation to Cause
25. Argument from Evidence to a Hypothesis
26. Abductive Reasoning
27. Argument from Position to Know
28. Argument from Expert Opinion
29. Argument from Waste

Later work [19] presented a compendium of ninety-six argumentation schemes, depending on how the subtypes are classified.

For example, argument from expert opinion is a common type of argument that we are often interested in for argumentation studies. It is made up of two distinctive premises and a conclusion. Basically it says: so-and-so is an expert, so-and-so says that some proposition is true, therefore (defeasibly) this proposition is true. Identifying this particular type of argument would seem to be simple. For example, the method could use keywords, like the word 'expert'. However, from experiences with teaching informal logic methods to students, there is a problem that occurs with some students who will immediately go to the Internet when asked to find examples of this kind of argument, and pick the first text they find containing the word 'expert'. Of course, many of these examples are not instances of argument from expert opinion. Keywords that occur in standardized forms of arguments, like the word 'expert', can be useful in helping a student to find examples of a specific type of argument. But they are crude tools, because their use without further refinement results in many errors.

This kind of work represents a more systematic continuation of the kind of practice that is carried out in teaching courses on argumentation or informal logic. Over many years of teaching courses of this type, I always used basically the same method of starting to teach the students through the use of examples. I searched through magazines, newspapers, and the Internet, or whatever other sources of material were available, to find interesting examples of arguments from expert opinion, ad hominem arguments, cases of equivocation, and so forth. From building up stocks of these cases and discussing and analyzing them with my classes, I started to build up accounts of each of the types of arguments, the kinds of premises they have, and the different varieties of them.

However, I didn't do this collecting in any systematic way. The examples I found initially came from the news magazine I usually read, or from the sections on informal logic in the many logic textbooks that use such examples. Eventually the wealth of experience that came from studying these examples led to the formulation of argumentation schemes, forms used to represent the basic structure of each type of argument. The schemes turned out to be very helpful as I continued to teach courses on argumentation, because they gave students some guideposts to use in their attempts to identify, analyze and evaluate arguments.

There were two kinds of assignments I used to typically give to my students in these courses. In the one type, I gave them each the same text of discourse containing an interesting argument, say a one page magazine opinion editorial. In the other type of assignment, I asked the students themselves to find an interesting example of one of the arguments we were concerned within the class, like argument from expert opinion, and analyze and evaluate their example. These tasks correspond to what the method to be built in this project is designed to help with. So it is easily seen how such a method would be helpful for teaching courses of this sort. It would also have a much wider use, however. For example, it would be an extremely powerful tool for researchers in fields like argumentation and informal logic. They could collect masses of interesting data on particular types of arguments that have long been studied in a more anecdotal way, and make the findings of the field of argumentation study much more powerful, because it is would then be based on documented data of a comprehensive sort.

Another example that illustrates how the project will work is the ad hominem type of argument. The way I defined this type of argument and crafted the argumentation schemes for it, there has to be more than just a personal attack. For something to be a

genuine ad hominem argument, four requirements have to be met. First, there have to be two arguers who are engaging in some sort of argumentation with each other. Second, one of the arguers has to have put forward an argument. Third, the other arguer has to be attacking the first party's argument. And fourth, the other arguer has to be using personal attack for this purpose. Very often I found that if I asked students to go and collect an interesting example of an ad hominem argument, they would find some instance of name-calling or personal attack, "Bob is a liar", and label that as an instance of an ad hominem argument. But if the instance of name-calling was not being used to attack somebody's argument, according to the argumentation scheme for the *ad hominem* argument, it should not correctly be so classified. Of course, one can debate the classification system, and there has been plenty of that going on in the field of informal logic, but to carry out a systematic study of any domain, one has to start with some initial hypotheses, definitions and classifications of the things being studied. Hence it is most useful, and in my opinion even necessary, to start with a well-defined set of argumentation schemes, even if the definitions of them are only regarded as tentative hypotheses that are subject to modification and revision as the project processes more and more examples of a given type of argument.

2. Introducing Argumentation Schemes

Argument from expert opinion is a subspecies of position to know reasoning, based on the assumption that the source is in a position to know because she is an expert. In trying to apply these schemes to real cases of argumentation, it can sometimes be easy to get them mixed up. Here is a typical example of argument from position to know.

If one is trying to find the best way to get to City Hall in an unfamiliar city, it may be helpful to ask a passer-by. If it looks like this passer-by is familiar with the city, and she says that City Hall is 12 blocks east, it could be reasonable to accept the conclusion that City Hall is 12 blocks east.

This form of reasoning is called position to know argumentation.
Where *a* is a source of information, the following argumentation scheme represents the form of position to know argumentation.

Scheme for Argument from Position to Know [19, 309]

MAJOR PREMISE: Source *a* is in position to know about things in a certain subject domain *S* containing proposition *A*.
MINOR PREMISE: *a* asserts that *A* is true (false).
CONCLUSION: *A* is true (false).

Such an argument can be reasonable in many instances, but it also defeasible. It can be critically questioned by raising doubts about the truth of either premise, or by asking whether *a* is an honest (trustworthy) source of information. The following critical questions match the scheme for the position to know argument.

Critical questions

CQ$_1$: Is *a* in position to know whether *A* is true (false)?
CQ$_2$: Is *a* an honest (trustworthy, reliable) source?
CQ$_3$: Did *a* assert that *A* is true (false)?

The second critical question concerns the credibility of the source. For example, a lawyer cross-examining a witness in a trial is allowed (within controlled limits) to raise critical questions about the character of the witness for honesty. If a witness has been known to lie in previous cases, a cross-examiner is allowed to ask such ad hominem questions, as an exception to the general rule against prejudicing the jury, by attacking the ethical character of a defendant.

Let us consider once again the case of asking the passerby where City Hall is located in a city one is not familiar with. Such a case is clearly an instance of position to know reasoning, but is it also an instance of the scheme for argument from expert opinion? Students in a critical thinking course are often inclined to think so, because it may seem to them reasonable to say that the passerby is being consulted as an expert on the city streets. After all, if she is very familiar with them, she might be said to have a kind of expert knowledge of them.

The argumentation scheme for argument from expert opinion is different from the one for argument from position to know, because it is required that the source who is in a position to know be an expert. For example, ballistics experts and DNA experts are often used to give expert testimony as evidence in trials, but they must qualify as experts. The basic version of this scheme for argument from expert opinion is given in [19, 310].

MAJOR PREMISE: Source *E* is an expert in subject domain *S* containing proposition *A*.
MINOR PREMISE: *E* asserts that proposition *A* is true (false).
CONCLUSION: *A* is true (false).

It is rarely wise to treat an expert as an infallible source of knowledge, and taking that approach makes argumentation susceptible to the fallacious misuse of argument from expert opinion. Generally, this form of argumentation is best treated as defeasible, subject to failure under critical questioning. The six basic critical questions matching the appeal to expert opinion [16, 223] are the following.

1. *Expertise Question*: How credible is *E* as an expert source?
2. *Field Question*: Is *E* an expert in the field *F* that *A* is in?
3. *Opinion Question*: What did *E* assert that implies *A*?
4. *Trustworthiness Question*: Is *E* personally reliable as a source?
5. *Consistency Question*: Is *A* consistent with what other experts assert?
6. *Backup Evidence Question*: Is *E*'s assertion based on evidence?

If the respondent asks any one of the six critical questions a burden of proof shifts back to the proponent's side to respond appropriately.

On this interpretation, part of the argument in the City Hall example would fit the major premise of the scheme for argument from expert opinion. But in the absence of further evidence, can it be correctly said that she is an expert? Unless she is a cartographer, or an expert on city planning, or has some qualification of that sort, she

would not qualify as an expert in the sense in which the term is used in law. In short, we can draw a distinction between having a working or practical knowledge of some area, and having expert knowledge of it.

It might be noted here as well that many arguments that occur in real argumentation texts, whether in law or everyday conversational argumentation, have implicit premises or conclusions [18, 2005]. Consider the example, "Jao lives in Lisbon and says the weather is fine there, therefore the weather is fine there". An implicit premise is that Jao is in a position to know about the weather in Lisbon, based on the explicit premise that he lives there. Another implicit premise is the defeasible conditional that if a person lives in a place, he is in a position to know about the weather there. The implicit conclusion of the argument is the statement that the weather is fine in Lisbon. At present such implicit premises and conclusions can only be found by having a human analyst dig them out as best explanations of the meaning of the text. It should be noted that argumentation schemes are very helpful for this purpose in many instances. In the Lisbon example, the scheme for argument from position to know can be applied to extract the missing premise and conclusion. We use an argument diagramming tool called Arauacaria, found at http://araucaria.computing.dundee.ac.uk/doku.php to insert the conclusion of the argument that the weather is fine in Lisbon and the premise that Jao is in a position to know about the weather in Lisbon. We can determine that there is an implicit premise in this argument, the propositions that Jao is in a position to know about the weather in Lisbon. We can find this missing a premise by selecting the argumentation scheme for 'argument from position to know' shown in figure 1.

Figure 1: Schemes Menu for Argument from Position to Know in Araucaria

On the left we see the argumentation scheme for argument from position to know. On the right we see the particular argument about the weather in Lisbon that fits the requirements of this scheme. We also see at the bottom of the menu some critical questions that can be used to respond to an argument fitting the scheme.

Now we can draw a diagram representing the premises and conclusion of the argument, and the arguments joining them. In figure 2 the missing premise, the statement that Jao is in a position to know about the weather in Lisbon is inserted in a darkened box with a broken border in the middle of the argument diagram.

Argument from Position to Know

The weather is fine in Lisbon.

Jao says the weather is fine in Lisbon.

Defeasible Modus Ponens

Jao is in a position to know about the weather in Lisbon.

Jao lives in Lisbon.

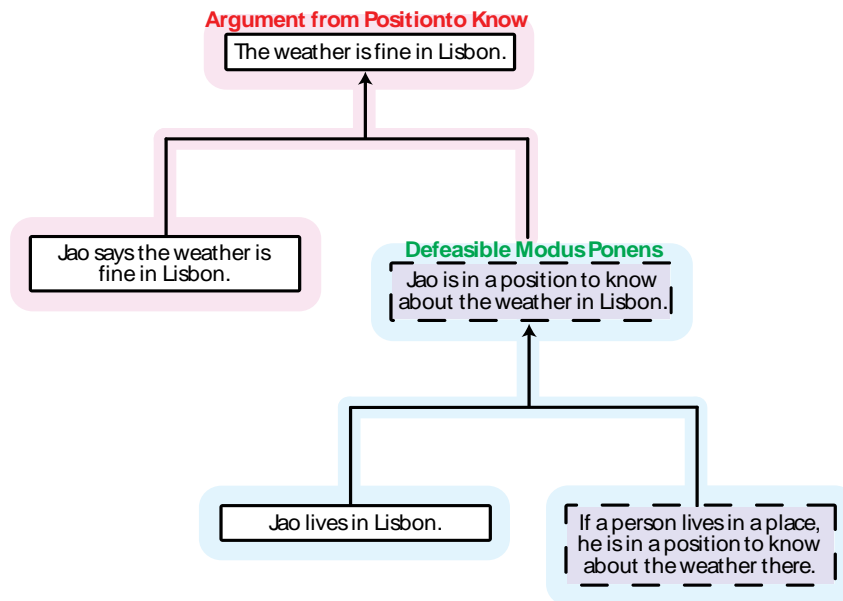If a person lives in a place, he is in a position to know about the weather there.

Figure 2: Argument Diagram for the Lisbon Example

As shown in figure 2, the missing premise that Jao is in a position to know about the weather in Lisbon is shown as being supported by two other premises. One is the explicit premise that Jao lives in Lisbon, and the other is the implicit premise that if a person lives in a place he is in a position to know about the weather there. This argument is shown as having the argumentation scheme for defeasible modus ponens. Modus ponens can take two forms, strict modus ponens and defeasible modus ponens. Defeasible modus ponens has a conditional premise that is open to exceptions. Defeasible modus ponens has the following form, where $A \Rightarrow B$ is the defeasible conditional: $A \Rightarrow B$; A; therefore B. For example, if something is a bird then generally, subject to exceptions, it flies; Tweety is a bird; therefore Tweety flies. If we find out that Tweety is a penguin, the original defeasible modus ponens argument defaults.

Argument from ignorance, or lack-of-evidence reasoning as it is often called, is another scheme that is so common and natural to use that it is hard to identify. We use it all the time, but are scarcely aware we are doing it. This scheme is hard for students to grasp at first and to identify in natural language discourse, because it is subtle, and because it involves negation. The scheme for argument from ignorance is based both what is known and what is not known to be true at some point in a sequence of argumentation [15, 254]. The major premise is a counterfactual.

MAJOR PREMISE: If A were true, A would be known to be true.
MINOR PREMISE: A is not known to be true.
CONCLUSION: A is false.

The major premise is based on the assumption that there has been a search through the knowledge base that would contain A that has supposedly been deep enough so that if A were there, it would be found. The critical questions include considerations of (1) how deep the search has been, and (2) how deep the search needs to be to prove the conclusion that A is false to the required standard of proof in the investigation. In typical instances of the argument from ignorance, the major premise of the argument is not explicitly stated. It has to be extracted from the text by applying the argumentation scheme. It is perhaps also for this reason that students tend to overlook this type of argument and have a hard time identifying it in a given text.

3. Schemes for Practical Reasoning and Arguments from Consequences

There are variants of the scheme for practical reasoning, but the simplest one (below) is called practical inference. In the scheme below, the first-person pronoun 'I' represents a rational agent of the kind described by Woodridge [20], an entity that has goals, some (though possibly incomplete) knowledge of its circumstances, and the capability of acting to alter those circumstances and to perceive (some of) the consequences of so acting. The simplest form of practical reasoning is the following fast and frugal heuristic for practical inference [19, 323].

MAJOR PREMISE: I have a goal G.
MINOR PREMISE: Carrying out this action A is a means to realize G.
CONCLUSION: Therefore, I ought (practically speaking) to carry out this action A.

Below is the set of critical questions matching the scheme for practical inference [19, 323].

CQ_1: What other goals do I have that should be considered that might conflict with G?
CQ_2: What alternative actions to my bringing about A that would also bring about G should be considered?
CQ_3: Among bringing about A and these alternative actions, which is arguably the most efficient?
CQ_4: What grounds are there for arguing that it is practically possible for me to bring about A?
CQ_5: What consequences of my bringing about A should also be taken into account?

The last critical question, CQ_5, is very often called the side effects question. It concerns potential negative consequences of a proposed course of actions. Just asking about consequences of a course of action being contemplated could be enough to cast an argument based on practical reasoning into doubt.

Another possibility is that an argument based on practical reasoning could be attacked by the respondent claiming that there are negative consequences of the proposed action. This move in argumentation is stronger than merely asking CQ$_5$ as it is an attempted rebuttal of the original argument. There is a specific argumentation scheme representing this type of argument. Argument from negative consequences cites the consequences of a proposed course of action as a reason against taking that course of action. This type of argument also has a positive form, in which positive consequences of an action are cited as a reason for carrying it out. These are the two basic argumentation schemes for arguments from consequences [19, 332], where *A* represents a state that could be brought about by an agent.

Argument from Positive Consequences

PREMISE: If *A* is brought about, good consequences will plausibly occur.
CONCLUSION: Therefore *A* should be brought about.

Argument from Negative Consequences

PREMISE: If *A* is brought about, then bad consequences will occur.
CONCLUSION: Therefore *A* should not be brought about.

Argumentation from consequences offers a reason to accept a proposal for action tentatively, subject to exceptions that may arise as new circumstances become known. An instance of argument from consequences can be stronger or weaker, depending on its initial plausibility and the critical questions that have been used to attack it.
    The scheme for argument from positive value [19, 321] takes the following form.

PREMISE 1: Value *V* is *positive* as judged by agent *A*.
PREMISE 2: If *V* is *positive*, it is a reason for *A* to commit to goal *G*.
CONCLUSION: *V* is a reason for *A* to commit to goal *G*.

The scheme for argument from negative value [19, 321] takes the following form.

PREMISE 1: Value *V* is *negative* as judged by agent *A*.
PREMISE 2: If *V* is *negative*, it is a reason for retracting commitment to goal *G*.
CONCLUSION: *V* is a reason for retracting commitment to goal *G*.

Argument from negative consequences is a form of rebuttal that cites the consequences of a proposed course of action as a reason against taking that course of action.
    Another type of argument widely used in the argumentation is the variant of practical reasoning called value-based practical reasoning [3, 2003]. The version of this scheme below is from (Walton, Reed and Macagno, 2008, 324).

PREMISE 1: I have a goal *G*.
PREMISE 2: *G* is supported by my set of values, *V*.
PREMISE 3: Bringing about *A* is necessary (or sufficient) for me to bring about *G*.

CONCLUSION: Therefore, I should (practically ought to) bring about *A*.

Another version of the scheme for value-based practical reasoning [2, 88] unpacks the notion of a goal into three elements: the state of affairs brought about by the action, the goal (the desired features in that state of affairs) and the value (the reason why those features are desirable).

In the current circumstances *R*
Action *A* should be performed
To bring about new circumstances *S*
Which will realize goal *G*
And promote value *V*

Note that value-based practical reasoning can be classified as a hybrid scheme that combines argument from values with practical reasoning.

In some cases, it may be hard to identify the type of an argument, because it is not obvious what scheme it fits, because some elements of the argument are not explicitly stated. However, in such cases, clues from the context of dialogue can help. Consider the following blood pressure dialogue [12, 3].

Proponent: Have you had your blood pressure checked?
Respondent: There is no need.
Proponent: Uncontrolled high blood pressure can lead to heart attack, heart failure, stroke or kidney failure.

The respondent's reply, 'There is no need', could be seen as a way of attacking the second premise of the simplest version of the scheme for argument from practical reasoning. He is saying, in effect, that having his blood pressure checked is not a necessary means to maintain his health. This move illustrates the type of rebuttal that is an attack on a premise of an argument. As noted above, the fifth critical question is associated with argument from negative consequence.

It may be possible to reconstruct the proponent's reaction at his first move by using this clue. It could be interpreted as an instance of the scheme for practical reasoning if we insert an implicit premise. [12] offers no further information about the context of the argumentation in the example, but it seems reasonable to presume that the proponent is concerned about the respondent's health. If so, one could insert as an implicit premise the statement that a goal for the respondent is his health. If this is a reasonable assumption, the proponent's argument could be reconstructed as follows

Implicit Premise: Your goal is to maintain your health.
Explicit Premise: Having your blood pressure checked is a necessary means to maintain your health.
Conclusion: You should have your blood pressure checked.

On this interpretation, the scheme for practical reasoning can help to reconstruct the argumentation sequence, as one can see as follows. The proponent made an argument

from practical reasoning. The respondent questioned the major premise of this practical argument. He doubts that having your blood pressure checked is necessary to maintain your health. The proponent than provided an argument to support the major premise of his practical argument. The above analysis of the example is meant to be simple, for purposes of illustration. A fuller analysis would show how another scheme, value-based practical reasoning, is involved, and also how it is the necessary condition variant of the scheme for practical reasoning that is involved.

4. The Sunk Costs Argument

Argument from waste is a kind of argument where one party thinking of discontinuing some course of action she has been engaging in for some time, and another party argues, "If you stop now, all your previous efforts will be wasted". It is also called the sunk costs argument in economics, where it has traditionally been regarded as a fallacy, even though more recently, it has been thought to be reasonable in many instances. It is typified by the following kind of case [17, 473].

Someone has invested a significant amount of money in a stock or business. Decreasing value and poor performance suggest it might be a good time to pull out and invest the remaining money elsewhere. But because the person has already invested so much in this venture, and would lose so much of it by pulling out now, she feels that she must stay with it rather than take the loss. To abandon the investment would be too much of a waste to bear, given all the money that has been sunk into it at this point. Reasoning on the basis of sunk costs, the person concludes that she must stay with this investment, even though she is convinced that the prospects for its rising in value are not good.

It is not hard to see why the sunk costs argument is often regarded as fallacious in economics and in business-decision-making generally, where an investor needs to think of the future, and should not be emotionally tied to previous commitments once circumstances change. However, there are other cases where the sunk costs argument can be reasonable [17], especially those where one's commitments to something one has put a lot of effort into are based on one's values. The sunk costs argument appears to be a species of practical reasoning that is also built on argument from consequences and argument from values. It is a composite argument built from these simpler schemes. So it is a classic case raising the problem of how this cluster of schemes should be structured in a classification system for argumentation schemes.

The sunk costs argument is a subtype of argument from negative consequences, as can be seen by putting it in this form.

If you stop doing what you are doing now, that would be a waste.
Waste is a bad thing (negative consequence).
Therefore you should not stop doing what you are doing now.

More precisely, the sunk costs argument has the following argumentation scheme, where *A* is an outcome of an action, argument from and *a* is an agent [19, 326].

PREMISE 1: If *a* stops trying to realize *A* now, all *a*'s previous efforts to
            realize *A* will be wasted.
PREMISE 2: If all *a*'s previous attempts to realize *A* are wasted, that would

be a bad thing.

CONCLUSION: Therefore, *a* ought not to stop trying to realize *A*.

In [17] it is shown that the sunk costs argument can a reasonable form of argument, but also that it is defeasible, and open to the following critical questions.

CQ_1: Is bringing about *A* possible?

CQ_2: If past losses cannot be recouped, should a reassessment of the cost and benefits of trying to bring about *A* from this point in time to be made?

A failure to address a critical question appropriately during changing circumstances, due to an attachment to previous commitments and efforts, is associated with fallacious instances of argument from sunk costs.

As indicated above, the sunk costs argument needs to be as built on argument from negative consequences and argument from negative values. This takes us to the problem of classifying schemes, by building classification trees showing one scheme is a subspecies of another. This problem will be taken up in section 6.

Another problem we now need to address is of a different sort, although it is related to classification as well. It has to do with recognizing an instance of a scheme in a natural language argument in a given text of discourse. In some instances, a given argument may look like it should be classified as an instance of the scheme for argument from waste, but there may be questions about this classification. An example can be found in an opinion article in the *Western Courier*, October 25, 2008.[1] The article advocates the use of embryonic stem cells for the advancement of medicine, and claims that the technology exists for deriving human embryonic stem cells without harming the embryo. The argument opposes the position of conservative groups, who are unwilling to support any kind of embryonic research, regardless of whether or not it destroys the embryo. One of the main arguments appears to be an instance of argument from waste, as indicated in the part of the article quoted below.

The Stem Cell Example

This [position] is shortsighted and stubborn. The fact is, fetuses are being aborted whether conservatives like it or not. Post-abortion, the embryos are literally being thrown away when they could be used in life-saving medical research. It has become a matter of religious and personal beliefs, and misguided ones at that. Lives could be saved and vastly improved if only scientists were allowed to use embryos that are otherwise being tossed in the garbage.

The argument in this article could be recast in a format that makes it appear to fit the scheme for argument from waste.

Premise 1: The embryos could be used in lifesaving medical research.
Implicit premise 1: Lifesaving medical research is a good thing.
Premise 2: The embryos are being thrown away.

---

[1] (www.westerncourier.com/news/2006/09/01/Opinion/**Stem-Cells**.Are.**Going.To.Waste**-2255161.shtml - 40k -

Implicit premise 2: Anything being thrown away that could be used is a waste [a bad thing].
Conclusion: The embryos should be used in medical research.

Putting the argument in this format makes it appear that it is an instance of argument from waste. But is it? This question is puzzling, and opinions on both sides can be found. The word 'waste' is used, and waste is taken to be a bad thing. Also, that some action is classified as a "waste" is taken as a reason against it. However, what appears to be missing is that in a proper argument from waste, as required by premise 1 of the scheme, the agent was making some previous efforts to do something, and if he stops now, his efforts will be wasted. In the stem cell example, there were no previous efforts of this sort. Instead, what is said to be a waste are the embryos that are "thrown away". Premise 2 of the scheme for argument from waste also requires that if previous attempts to realize something are wasted, that would be a bad thing. There seems to be nothing fitting this premise in the stem cells example. Nobody was doing anything with the stem cells previously. No effort or commitment was being put into doing something with them.

The problem with this kind of case is one of matching the premises of the scheme with the premises of the argument given in the text of discourse. If a required premise of the scheme is not found in the argument in the discourse, we need to look and see if there is evidence that it is an implicit premise. If there is no such evidence we need to conclude that the given argument is not an instance of this particular scheme. In the present example, we need to conclude that the given argument is not an instance of the scheme for the sunk costs argument.

5. Slippery Slope Arguments

Finally we need to consider the scheme for the slippery slope argument, a highly complex form of argument composed of the other simpler schemes we have so far studied. It is a common problem in teaching critical thinking skills to students that once they are taught a structure for the slippery slope argument, they tend to apply it to cases where the evidence does not really justify classifying it under this category. For example, they typically find cases of argumentation from negative consequences and leap to the conclusion that it must be a slippery slope argument because some bad outcome is being cited as a reason for not carrying out a particular course of action. Strictly speaking, however, to be a slippery slope, an argument has to meet a number of requirements.

Scheme for the Slippery Slope Argument [19, 339]

FIRST STEP PREMISE: A_0 is up for consideration as a proposal that seems initially like something that should be brought about.
RECURSIVE PREMISE: Bringing up A_0 would plausibly lead (in the given circumstances, as far as we know) to A_1, which would in turn plausibly lead to A_2, and so forth, through the sequence $A_2, \ldots A_n$.
BAD OUTCOME PREMISE: A_n is a horrible (disastrous, bad) outcome.
CONCLUSION: A_0 should not be brought about.

It is an important requirement for this scheme that the recursive premise be present. Without that premise, the argument clearly is simply an instance of argument from negative consequences. It is the presence of the recursive premise that enables us to distinguish any given case whether the argument is a slippery slope argument as well as being in argument from negative consequences.

The following example is a genuine slippery slope argument. It concerned the burning of an American flag by Gregory Lee Johnson during a political demonstration in Dallas to protest policies of the Reagan administration. Johnson was convicted of "desecration of a venerated object", but the Texas Court of Criminal Appeals reversed the ruling, arguing that Johnson's act was "expressive conduct", protected by the First Amendment. This decision was reaffirmed by the Supreme Court
in the case of Texas v. Johnson (1989 WL 65231(U.S.), 57 U.S.L.W. 4770). In delivering the opinion of the Court, Justice William Brennan cited the precedent case of Schacht v. United States, where it was ruled that an actor could wear a uniform of one of the U.S. armed forces while portraying someone who discredited that armed force by opposing the war in Vietnam.

We perceive no basis on which to hold that the principle underlying our decision in Schacht does not apply to this case. To conclude that the Government may permit designated symbols to be used to communicate only a limited set of messages would be to enter territory having no discernible or defensible boundaries. Could the Government, on this theory, prohibit the burning of state flags? Of copies of the Presidential seal? Of the Constitution? In evaluating these choices under the First Amendment, how would we decide which symbols were sufficiently special to warrant this unique status? To do so, we would be forced to consult our own political preferences, and impose them on the citizenry, in the very way that the First Amendment forbids us to do.

The argument in this text can be identified as an instance of the slippery slope argument. A first step is said to lead to a series of unclear decisions (whether to prohibit the burning of state flags, copies of the Presidential seal, the Constitution, and so forth), which would in turn lead to the outcome of individuals imposing their own political preferences on the citizens. This is said to be an intolerable outcome in a free country, a violation of the First Amendment. In this instance, clearly the recursive premise is present.

The problem with many examples is that the argument does appear to be of the slippery slope type, but the series of intervening steps required to get from the premises to the conclusion is not made explicit. The following excerpt from a letter written by Richard Nixon in *The New York Times* on October 29, 1965, has been taken as an example of a fallacious slippery slope argument in logic textbooks (Walton, 1992, 97), but it is hard to tell whether it really fits this scheme for slippery slope argument. Nixon's letter warned about consequences of the fall of Vietnam in the following terms.

. . . would mean ultimately the destruction of freedom of speech for all men for all time not only in Asia but the United States as well. . . . We must never forget that if the war in Vietnam is lost . . . the right of free speech will be extinguished throughout the world.

This argument certainly looks like a classic case of the slippery slope argument, but where is the recursive premise? The intervening steps are missing. Presumably what Nixon was claiming is that the fall of Vietnam would lead to the fall of other neighboring countries to communism, and these events in turn would cause a chain reaction with the

final disastrous outcome that the whole world is taken over by undemocratic countries. The problem is that Nixon did not fill in all these intervening steps, and so how can we prove that the recursive premise requirement is really meant by the argument as stated in the example above? One option that needs to be looked at is whether these intervening claims can be taken to be implicit premises. In other words, is the argument an enthymeme? There is evidence for this contention, and thus by marshalling the textual evidence, a case can be made that the argument should properly be classified as a slippery slope. However, the contention that this argument fits the scheme for the slippery slope argument needs to be argued for. If it cannot be sustained by the marshalling of the textual and contextual evidence in the case, the argument should only be classified as an argument from negative consequences that is not also a slippery slope argument.

The slippery slope argument does occur in everyday and legal argumentation, but it is not nearly as common as other schemes mentioned above, like argument from negative consequences and practical reasoning. As shown by the example above, it is a substantial task to properly identify an argument in a given case as fitting the slippery slope scheme, because the scheme is so complex, with so many prerequisites, and because it is a composite, made up of other simpler schemes. It is argued in [14] that the slippery slope scheme can be analyzed as a complex chain of subargments each having the DMP structure, but there is insufficient space to discuss this interesting analysis here.

6. Classification of Schemes

The project of automatic identification of arguments in a text using schemes would greatly benefit from a classification system showing which schemes are subschemes of others. The subject of classification schemes is a topic for another paper, but it will help to make a few comments on this related project here, since it is so obviously important in the schemes and examples studied above. So far there is no generally accepted system of classification for argumentation schemes. Walton, Reed and Macagno [19, 349-350] surveyed several different approaches, and concluded that it appears that at the present state of development of schemes, the general system summarized below is the easiest to apply. In it, there are three main categories, and various schemes under each one.

REASONING

1. *Deductive Reasoning*
   Deductive *Modus Ponens*
   Disjunctive Syllogism
   Hypothetical Syllogism
   Reductio Ad Absurdum
   Etc.

2. *Inductive Reasoning*
   Argument from a Random Sample to a Population
   Etc.

3. *Practical Reasoning*

Argument from Consequences
Argument from Alternatives
Argument from Waste
Argument from Sunk Costs
Argument from Threat
Argument from Danger Appeal

4. *Abductive Reasoning*
   Argument from Sign
   Argument from Evidence to a Hypothesis

5. *Causal Reasoning*
   Argument from Cause to Effect
   Argument from Correlation to Cause
   Causal Slippery Slope Argument
(For details see chapter 5 of [19] on causal argumentation).

SOURCE-BASED ARGUMENTS

1. *Arguments from Position to Know*
   Argument from Position to Know
   Argument from Witness Testimony
   Argument from Expert Opinion
   Argument from Ignorance

2. *Arguments from Commitment*
   Argument from Inconsistent Commitment

3. *Arguments Attacking Personal Credibility*
   Arguments from Allegation of Bias
   Poisoning the Well by Alleging Group Bias
   Ad Hominem Arguments
   Etc.

4. *Arguments from General Acceptance*
   Argument from Generally Accepted Opinion
   Argument from General Practice
   Etc.

APPLYING RULES TO CASES

1. *Arguments Based on Cases*
   Argument from Example
   Argument from Analogy
   Argument from Precedent
2. *Defeasible Rule-Based Arguments*

Argument from an Established Rule
Argument from an Exceptional Case
Argument from Plea for Excuse
3. *Verbal Classification Arguments*
Argument from Verbal Classification
Argument from Vagueness of a Verbal Classification
4. *Chained Arguments Connecting Rules and Cases*
Argument from Gradualism
Precedent Slippery Slope Argument
Slippery Slope Argument

This classification scheme is very helpful for the purpose of orienting students taking an informal logic course, because it helps group some of the most commonly used schemes into categories. But as shown in the instances of the slippery slope argument and argument from sunk costs, the relationship of these more complex schemes to the simpler schemes that compose them, like practical reasoning, argument from consequences and argument from values, requires a deeper analysis.

Prakken [10] has given another example of how to schemes are structurally related in an interesting way. He studied the structural relationship between argument from expert opinion and argument from position to know, and showed that the former scheme can be classified as a special instance of the latter scheme. He also showed how argument from evidence to a hypothesis can be analyzed in a manner showing that it is a species of abductive reasoning, often called inference to the best explanation. These findings confirm the hypothesis that many of the most common schemes have an interlocking relationship with other schemes so that one scheme can be classified as a subspecies of another, but only in a complex manner. This complex manner needs to take into account structural relationships between the schemes.

What has been shown here is very important for not only developing a precise system of classification of schemes, but also for the overarching project of developing a system for argument mining. What has been shown is the there are some simple and basic schemes, and there are some highly complex schemes that are built up as complexes from the simpler schemes. Among the most important simple schemes are practical reasoning, argument from position to know, argument from commitment, and argument from values. Argument from consequences is also a simple scheme, but it has an interesting relationship to the scheme for practical reasoning. Argument from negative consequences corresponds to one of the critical questions matching the scheme for practical reasoning.

There is not sufficient space in this paper to take up the problem of how to provide a deeper classification of argumentation schemes of a kind that would be useful for argument mining and other informal logic tasks. It must be left as a problem for future research. In the next section, is shown that there has been some groundbreaking work in artificial intelligence on taking the first steps in building systems of argument mining. It is very useful for those of us in the informal logic field to know about this research.

7. Research on Argument Mining in Artificial Intelligence

It is encouraging that there are already some systems applying argumentation schemes to legal texts that have been implemented and the results of this experimental work are very interesting so far [8, 27]. Discourse theories assume that the structure of a text is that of a graph or a tree, and that the elementary units of complex text structures are non-overlapping spans of text. [5] conducted experiments directed towards the ultimate aim of developing methods for automatically classifying arguments in legal texts in order to make it possible to conveniently access and search types of arguments in such texts. They build on recent work in legal argumentation theory as well as rhetorical structure theory. They look for prominent indicators of rhetorical structure expressed by conjunctions and certain kinds of adverbial groups [5, 226]. They identify words, pairs of successive words, sequences of three successive words, adverbs, verbs and modal auxiliary verbs. Rhetorical structure theory defines twenty-three rhetorical relations that can hold between spans of a text. Most hold between two text spans called a nucleus, the unit most central to the writer's purpose, and a satellite, which stands in a relation to the nucleus. For example, the evidence relation links a nucleus like 'Bob shot Ed' and a satellite like 'Bob's fingerprints were found on the gun'. Their experiments offer an initial assessment of types of features that play a role in identifying legal arguments and single sentences. In future work, they also hope to focus on the classification of different types of arguments.

This work has been applied to legal argumentative texts [6]. In this research [8], the sentences are classified according to argumentation schemes, and the aim is to build a system for automatic detection and classification of arguments in legal cases [7]. The project has built a corpus from texts of the European Court of Human Rights that was annotated by three annotators under supervision of a legal expert [8]. This task was made easier by the fact that the court documents that provided their corpus was already classified using subheadings into different parts of the text that had different functions. So, for example, there is one section of the text where the arguments of the judges are presented. Using a limited number of argumentation schemes, for example the 26 or so identified in [15] would be a way to start identifying the different types of arguments. This research opens up opportunities for applying artificial intelligence research to informal logic.

It is interesting to note that there were no identifications of instances of argument from ignorance in the corpus, and very few instances of argument from commitment were identified. Practical reasoning was not used as a scheme in this study. 80 instances of argument from position to know were found. 2099 instances of circumstantial argument against the person were found. 10744 instances of argument from evidence to a hypothesis were found. 2385 instances of argument from expert opinion were found. 12229 instances of argument from precedent were found. 1772 instances of arguments that fitted no scheme were found.

These results are interesting, but Mochales and Leven [8, 27] noted a number of problems. To improve the usefulness of systems for automated argument text mining, several research topics are acutely in need of exploration. These observations suggest that what is needed, in addition to the more precise definitions of the schemes themselves for use in automated argument detection [11], is the provision of additional criteria which can be of assistance in determining whether a scheme applies to a given argument in a text of discourse or not in problem cases. Sources for collecting such criteria can already be found in work in AI [9] and in argumentation theory [13]. The blood pressure example

showed that the context of the dialogue in a case needs to be taken into account, as well as the indicator words, in the task of detecting a scheme in discourse. The clue to determining that practical reasoning was the scheme fitting the argument in this case was the critical questioning matching the proponent's use of practical reasoning.

The Amsterdam School has been conducting research for some time on the task of identifying arguments in a text of discourse using so-called argumentative indicators like 'thus', 'therefore' and 'because' [13]. A large part of this research has focused on the task of distinguishing between an item in a text of discourse that may properly be taken to represent an argument, as opposed to some other speech act, like the putting forward an explanation or the making of a statement. Only a few argumentation schemes have been studied so far. These include argument from analogy, argument from sign and causal arguments.

Another approach [21] reconstructs legal case-based reasoning in terms of argumentation schemes. This approach uses a set of cases, factors, and comparisons between cases to instantiate argumentation schemes from which justifications for an outcome of the case at issue can be derived. These include argument from precedent and argument from analogy. Cases have a plaintiff, a defendant, a set of factors present in the case, and an outcome for the plaintiff or defendant [21, 139]. They identify and define what they call the main scheme in a case, including its premises and conclusion [21, 143]. One premise of this main scheme is called the factors preference premise, which states that one factor was preferred to another in a previous case decided in the plaintiff's favor. They then introduce a new argumentation scheme they call the preference-from-precedent scheme, that is used to support the factors preference premise of the main scheme. In the general literature on argumentation schemes, this legal scheme would be at a particular species of the scheme called argument from precedent [19, 344]. They identify other schemes as well, showing how arguments fitting schemes on both sides can be used to support or attack other arguments used in the case at issue. There are some features of their approach that are especially significant. They use an applicability assumption which arises because there might be a number of reasons why an argument put forward in a case is not a suitable precedent for that case. They also distinguish between three different kinds of premises in schemes, called ordinary premises, assumptions and exceptions [4, 2007]. This approach is especially significant because it shows how schemes can be used within the framework of legal case-based reasoning (Ashley, 2006), and especially because it shows how factors can be used to define legal schemes and apply them to argumentation a legal case.

Rahwan et al. [11] have advanced research on the automated identification of particular schemes by building the first ontology of argumentation schemes in description logic, showing how description logic inference techniques can be used to reason about automatic argument classification. An OWL-based system is implemented for argumentation support on the semantic web. At the highest level, three concepts are identified, called statements, schemes describing arguments made up of statements, and authors of statements. Different species of scheme are identified, including rule schemes, which describe the class of arguments, conflicts schemes and preference schemes. The schemes are classified by classifying their components: their ordinary premises, assumptions, exceptions and conclusions. Statements may be classified as declarative or imperative. For example, in the scheme for argument from position to know (see section

1 above), the class of statement *PositiontoHaveKnowledgeStmnt* is defined as a species of declarative statement associated with the property *formDescription*, 'agent *a* is in a position to know whether statement *A* is true or false', that describes its typical form [4, 8]. Using these categories, it is possible to fully describe a scheme, like the scheme for argument from position to know, by stating the necessary and sufficient conditions for an instance to be classified as falling under this type. Special types of schemes called conflict schemes are identified. The method of identifying schemes is implemented in a web-based system called *Avicenna* [4, 11-13]. A user can search arguments on the basis of keywords, structural features and other properties [4, 12].

It is clear even from this brief description of current work on argument mining in the field of computing, that this technical initiative would benefit greatly from more work on refining argumentation schemes. It is also clear that, even though this technical work is only a first step towards the development of useful argument mining technology, there are already interesting implications on how the methods currently being used can be adapted to the needs of informal logic.

8. Argument Mining as an Informal Logic Method

How could we use these results to develop argument mining methods that could be used to help students of informal logic identify arguments - for example, arguments of the kind they encounter in natural language texts in newspapers, magazines or on the Internet, and want to analyze using standard methods of informal logic? We already have some methods helpful for this sort of task. They include, for example, the use of argument mapping tools to identify the premises and conclusions of arguments, and to show how one argument is connected to another in a chain of argumentation. A second, and more ambitious project, is to develop an automated argumentation tool for argument mining. The idea is that this tool could go onto the Internet and collect arguments of specifically designated types, like argument from expert opinion for example. Of course, these two tasks are connected, because the most powerful method would likely turn out to be to combine both tasks by having trained human users apply the automated tool to identify arguments on a tentative basis in the text of discourse, and then correct the errors made by the automated tool.

There are six distinct tasks in this endeavor. The first task is the identification of arguments in a text of discourse, as opposed to other entities, like statements, questions or explanations. Carrying out this task requires some definition of what an argument is, as opposed to the speech acts that can often be confused with putting forward an argument, like offering an explanation. Part of this task is the identification of broad types of argument, like deductive and inductive arguments, as opposed to the third category sometimes called plausible arguments. The second task is the identification of specific, known argumentation schemes. The principal way of recognizing a particular argumentation scheme is to be able to identify the premises and the conclusion that make up that scheme. What is required is a parser that can recognize that not only the individual units of speech in one of these premises or conclusions, like nouns and verbs, but also recognize particular nouns can occur in a scheme, like 'expert', or particular phrases like 'position to know'. The third task is the deeper classification of argumentation schemes. The fourth task is the more precise formulation of schemes. This

can be carried out a number of ways. One way is to formulate schemes that can be applied in particular field. For example, the scheme for argument from expert opinion needs to be formulated in a more precise way in law than it is for purposes of analyzing ordinary conversational argumentation, because specific criteria for argumentation from expert opinion as a kind of evidence have already been established in law through legal precedents and court judgments. The fifth task is to develop criteria to enable the differentiation between schemes that appear similar to each other or closely related to each other. The sixth task is to develop techniques for minimizing errors in the identification of schemes in natural language text of discourse. As shown above, in some cases it is easy to mix up one scheme with another. Part of the task here is to develop a corpus of borderline problem cases of this sort, and work on criteria that can be used to solve the problem. An important part of this fifth task is to develop a deeper classification system for argumentation schemes.

There is much work to be done before any useful system of argument identification based on argument mining can be implemented in an informal logic setting. What is needed is to encourage those in the field of informal logic to carry out research projects on the subject. An initial problem for anyone setting up this kind of research project is to decide what kind of natural language texts of discourse should be used as the database. Textbooks in informal logic often take their examples from magazine and newspaper articles, but they sometimes include examples of legal arguments as well. One project would be to take a particular newsmagazine and try to identify instances of arguments found in it, as well as trying to identify the type of argument. A second project would be to use examples of legal argumentation of some sort. A third project would be to use the database of arguments in Debatepedia, or some similar online source that contains lots of interesting arguments pro and contra on controversial issues at any given time.

9. Conclusions

The work using schemes for argument mining in legal discourse suggests that in addition to the schemes themselves, additional information that cites specific requirements an argument has to meet to qualify as instance of a particular scheme would be extremely useful. For example, to help tell whether an argument in a text should best be classified under the heading of argument from position to know or argument from expert opinion, some requirements telling the argument annotator what kind of source qualifies as an expert source, like those discussed in section 1, would be useful. The way to build up such additional resources is to better integrate theoretical research on schemes with the work of testing its application to texts of discourse.

Even though the research work in artificial intelligence offers grounds for optimism about the feasibility of the project of automated argument detection using schemes, the problem cases posed above concerning argument from position to know, argument from expert opinion, the sunk costs argument and the slippery slope argument are reasons for concern. It needs to be noted that this task is related to another problem, the problem of enthymemes, or arguments found in a natural language text of discourse that have implicit premises or conclusions. As we saw in tackling the problem of differentiating between slippery slope arguments and arguments from negative consequences, much often depends on implicit premises, like the recursive premise, which the classification of

the argument requires, but that were not explicitly stated in the given text of discourse. This kind of problem is a central one for argumentation studies. Would it be possible to build an automated system that could detect enthymemes and fill in the missing premises or conclusions so that an analysis of the argument with its missing premises indicated could be provided by an argument visualization tool? The short answer is that it might be a lot more difficult to build such a useful tool of this kind than one might initially think [18], but the employment of schemes will very definitely be helpful as part of the tool. So the project is worth doing for purposes of informal logic, and is closely related to the underlying problem of developing a deeper classification system for argumentation schemes.

Acknowledgments

References

[1] Kevin Ashley, 2006, "Case-Based Reasoning", *Information Technology and Lawyers*, ed. A. R. Lodder and A. Oskamp, Dordrecht, Springer, 2006, 23-60.

[2] Katie Atkinson, Trevor Bench-Capon and Peter McBurney, 2004, "Justifying Practical Reasoning", in Floriana Grasso, Chris Reed and Giovanni Carenini (eds.), *Proceedings of the Fourth International Workshop on Computational Models of Natural Argument* (CMNA 2004), Valencia: Spain, 87-90.

[3] Trevor Bench-Capon, 2003, "Persuasion in Practical Argument Using Value-based Argumentation Frameworks", *Journal of Logic and Computation*, 13, 429-448.

[4] Thomas F. Gordon, Henry Prakken and Douglas Walton, 2007, "The Carneades Model of Argument and Burden of Proof", *Artificial Intelligence*, 171, 875-896.

[5] Marie-FrancineMoens, Raquel Mochales Palau, Erik Boiy and Chris Reed, 2007, "Automatic Detection of Arguments in Legal Texts", *Proceedings of the International Conference on AI and Law* (ICAIL 2007), Stanford, California, 2007, 225-230.

[6] Raquel Mochales Palau and Marie-Francine Moens, 2007, "Study on Sentence Relations in the Automatic Detection of Argumentation in Legal Cases", in Arno R. Lodder and Laurens Mommers (eds.), *Legal Knowledge and Information Systems: JURIX 2007, The Twentieth International Conference*, Amsterdam: IOS Press, 89-98.

[7] Raquel Mochales and Marie-Francine Moens, 2008, "Study on the Structure of Argumentation in Case Law", in Enrico Francesconi, Giovanni Sartor and Daniela Tiscornia (eds.), *Legal Knowledge and Information Systems: JURIX 2008, The Twenty-First International Conference*, Amsterdam: IOS Press, 11-20.

[8] Raquel Mochales and Aagje Leven, 2009, "Creating an Argument Corpus: Do Theories Apply to Real Arguments?", *Proceedings of the 12[th] International Conference on Artificial Intelligence and Law*, New York: Association for Computing Machinery, Inc., 2009, 21- 30.

[9] Raquel Mochales Palau and Marie-Francine Moens, 2009, "Argumentation Mining: The Detection, Classification and Structure of Arguments in Text", *Proceedings of the 12[th] International Conference on Artificial Intelligence and Law*, New York: Association for Computing Machinery, Inc., 98-107.

[10] Henry Prakken, 2010, "On the Nature of Argument Schemes", in Chris Reed and Christopher W. Tindale (eds.), *Dialectics, Dialogue and Argumentation: an Examination of Douglas Walton's Theory of Reasoning and Argument*, London: College Publications, 2010,167 -185

[11] Iyad Rahwan, Bita Banihashemi, Chris Reed and Douglas Walton, 2011, "Avicenna: Argumentation Support on the Semantic Web", *Knowledge Engineering Review*, to appear.

[12] Angelo Restificar, Syed S. Ali and Susan W. McRoy, 1999, "ARGUER: Using Argument Schemas for Argument Detection and Rebuttal in Dialogs", in Judy Kay (ed.), *UMP99: International Conference on User Modeling*, New York: Springer-Wien, 315-317.

[13] Frans H. van Eemeren, Peter Houtlosser and Francisca Snoeck Henkemans, 2007, *Argumentative Indicators in Discourse*, Dordrecht: Springer.

[14] Douglas Walton, 1992, *Slippery Slope Arguments*, Oxford: Oxford University Press, 1992.

[15] Douglas Walton, 1996. *Argumentation Schemes for Presumptive Reasoning*, Mahwah, New Jersey: Erlbaum.

[16] Douglas Walton, 1997, *Appeal to Expert Opinion*, University Park, Pennsylvania: Penn State Press.

[17] Douglas Walton, 2002, "The Sunk Costs Fallacy or Argument from Waste", *Argumentation*, 16, 473-503.

[18] Douglas Walton and Chris Reed, "Argumentation Schemes and Enthymemes", 2005, *Synthese: An International Journal for Epistemology, Methodology and Philosophy of Science*, 145, 339-370.

[19] Douglas Walton, Chris Reed and Fabrizio Macagno, 2008, *Argumentation Schemes*, Cambridge: Cambridge University Press.

[20] Michael Wooldridge, 2000, *Reasoning about Rational Agents*, Cambridge, Mass.: The MIT Press.

[21] Adam Wyner and Trevor Bench-Capon, 2007, "Argument Schemes for Legal Case-Based Reasoning", in Arno R. Lodder and Laurens Mommers (eds.), *Legal Knowledge and Information Systems: JURIX 2007, The Twentieth International Conference*, Amsterdam: IOS Press, 139-149.