

# Analyzing Open Source License Compatibility Issues with Carneades

Thomas F. Gordon  
Fraunhofer FOKUS  
Berlin, Germany

thomas.gordon@fokus.fraunhofer.de

## ABSTRACT

The Carneades software system provides support for constructing, evaluating and visualizing arguments, using formal representations of facts, concepts, defeasible rules and argumentation schemes. This paper illustrates features of Carneades with a prototype legal application for analyzing open source software license compatibility issues in particular cases. The Carneades system provides a unique combination of features that make to our knowledge applications of this kind possible for the first time.

## Keywords

legal knowledge representation, computational models of argument, copyright law

## 1. INTRODUCTION

This paper illustrates features of the Carneades argumentation system [14] with examples from a prototype legal application for analyzing open source license compatibility issues [15]. As Bing [3], Fiedler [7], McCarty [17] and many others have noted, legal argumentation is “not primarily deductive, but rather a modeling process of shaping an understanding of the facts, based on evidence, and an interpretation of the legal sources, to construct a theory for some legal conclusion” [3]. The parties in a legal dispute construct competing theories and argue about their relative merits. Carneades is designed to support all the steps in this process of theory construction, argumentation and evaluation.<sup>1</sup>

The Carneades software system is based on a well-founded formal model of structured argumentation with support for proof burdens and standards [9,10], now called Carneades Argument

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*ICAIL '11*, June 6-10, 2011, Pittsburgh, PA  
Copyright 2010 ACM 978-1-4503-0755-0/11/06...\$10.00.

Evaluation Structures (CAES). It has been formally proven that the Carneades model of argument is a specialization of both Prakken’s ASPIC+ model of structured argumentation [8] and Brewka’s Abstract Dialectical Frameworks [5] and thus an instantiation of Dung’s Abstract Argumentation Framework [4,8]. Carneades has also been shown by Governatori to be closely related to Defeasible Logic [16]. A formal model of abduction in Carneades argument evaluation structures has been developed [2], which is useful for identifying relevant issues and computing minimal sets of statements, called positions, which, if proven, would make some goal statement acceptable (in) or not acceptable (out) in a stage of a dialogue.

Building on this formal foundation, the Carneades software provides a number of tools for interactively constructing, evaluating and visualizing arguments, as well as computing positions. Arguments are constructed using formalizations of facts, concepts, defeasible rules and argumentation schemes [11,12]. Facts and concepts are represented using the Web Ontology Language (OWL), an XML schema for description logic [1], a subset of first-order logic and thus with a monotonic (strict) entailment relation. Legal rules and argumentation schemes [18] are both modeled as defeasible inference rules, represented in the Legal Knowledge Interchange Format (LKIF) [6]. The rules of alternative, competing theories of the law can be included in a single model.

A combination of forwards and backwards reasoning is used to construct arguments: a description logic reasoner constructs the deductive closure of the concepts and facts in a forwards manner; the Carneades rules engine uses backwards reasoning to apply the defeasible inference rules in a goal-directed and stratified way to the deductive closure of the description logic theory of facts and concepts. The LKIF rule language has been extended to provide a way to declare the domain of variables using predicates defined in OWL, similar to the way variables are typed in programming languages. These domain declarations provide important control information that enables the rule engine to iterate over instances of the domains to more efficiently instantiate the rules.

In addition to the arguments constructed automatically from a knowledge-base of facts, concepts and rules, arguments can manually entered into the system by the user. These arguments can be completely ad hoc or instantiations of argumentation schemes. The Carneades system currently includes a library of

<sup>1</sup> <http://carneades.berlios.de>

about 20 of Walton’s most important argumentation schemes along with a software assistant which steps the user through the process of selecting and instantiating schemes.

As the arguments are constructed and edited, they are visualized in an argument map [13]. The graphical user interface, called the Carneades Editor, supports argument evaluation by providing tools to accept and reject statements, assign proof standards and weigh arguments. After every modification, the underlying computational model of argument is used to update and visualize the acceptability status of statements in the map. The differential legal effects of competing theories can be analyzed by assuming their rules to be valid and then checking how this effects the acceptability of issues of interest in the argument map. Moreover Carneades provides a “find positions” assistant which can be used to abduce theories with desired legal effects.

The rest of this paper illustrates features of Carneades with examples from the prototype legal application for analyzing open source license compatibility issues. We start with examples from a simple OWL ontology for describing software licenses and use and derivation relationships between works of software. Next we show how to use the ontology to model the facts of a case. We then show how to model some rules of copyright law in LKIF, focusing on the issue of whether linking to a software library produces a derivative work. These models are then used to illustrate how Carneades can be used to construct, evaluate and visualize arguments about whether or not the project may publish its software using a particular open source license, i.e. whether a preferred license is compatible with the licenses of the software used by the project. We conclude by showing how abduction can be used to find positions that are helpful for proving, or disproving, depending on one’s standpoint and interests, that a license is compatible.

## 2. Concepts and Facts

Carneades uses the Web Ontology Language (OWL), a World Wide Web standard XML schema for representing and interchanging description logic knowledge bases. These knowledge bases have two parts, for concepts (TBox) and facts (ABox). The top-level concepts, called “classes” in OWL, for our application are:

- CopyrightLicense. Individual licenses, with which a particular legal entity, the licensor, grants rights to another legal entity, the licensee.
- CopyrightLicenseTemplate. Open source license templates, such as the GPL or BSD. A particular copyright license can be an instance of such a template.
- LegalEntity. Legal persons, such as humans, corporations and associations.
- LicenseTerm. The rights granted by a license and the conditions of the license, limiting the rights granted.
- Work. Various kinds of intellectual products protected by copyright, including software.

The Work class is for all works protectable by copyright. There is a SoftwareEntity subclass of Work, intended to cover all kinds of software artifacts, including not only source and object code, but also more abstract entities such as APIs and specifications.

The main property of software entities of interest for license compatibility issues is the isDerivedFrom property, expressing that one entity has been derived from another. This legal issue depends on the jurisdiction and the interpretation of the governing law by the courts. The ontology includes properties for representing various ways that software can use other software, such as compiledBy and linksTo. These properties are not from the domain of copyright law, but rather from the domain of software engineering.

In legal terms, the model of software concepts provides the means to represent the *material facts* of a software licensing case. The legal question is whether a particular use of software, such as linking, is sufficient to create a derivative work. In legal jargon, the question is whether a material fact, linking, can be *subsumed* under the legal concept of a derivative work. None of these use relations has been defined in the ontology to be a subproperty of isDerivedFrom, because these legal issues have not been resolved and we want to leave room to argue about them.

In addition to these use relations, the ontology includes properties for representing the compatibility of license templates and for representing the licenses which have been issued for particular works.

The software ontology was used to model an example software project, roughly based on the current version of the Carneades system. To illustrate, here is the model of some facts about the Carneades inference engine:

```
CarneadesEngine
  type: SoftwareLibrary
  implementedIn: Clojure
  compiledBy: ClojureCompiler
  linksStaticallyTo: ClojureLib
  linksStaticallyTo: Pellet
```

One of the libraries used by the Carneades engine is the Clojure library (ClojureLib) which is licensed using the Eclipse Public License (EPL) template, named EPL\_Template in the model. In the model, the ClojureLib does not have the EPL\_Template as its license, but rather has its own license, named ClojureLicense, that is an instance of the EPL template. This way of modeling licenses allows each instance of a template license to have its own licensor and licensee. This is necessary since the license and the template used to create the license are two different objects with different properties. For example, the copyright owners of the software and the template license are typically different persons.

Another library linked to by the Carneades engine uses the GNU AGPL license template. Since the EPL and AGPL are incompatible reciprocal licenses, a central question will be whether the Carneades engine can be linked to both of these libraries. This will depend on whether the linking of software to a library causes the software to be a derivative work of the library. We will return to this question after showing, in the next section, how alternative interpretations of copyright law can be modeled using LKIF rules.

## 3. RULES

Description logic (DL) is semantically a decidable subset of first-order logic. This means that the inferences of description logic reasoners are *strict*: if the axioms of a DL knowledge base are true in some domain, then all of the inferences made by a (correctly implemented) DL reasoner are necessarily also true, without

exception. While DL is very powerful and useful, monotonic logics are not sufficient for modeling legal rules, such as the rules of copyright law, in a maintainable and verifiable way, isomorphic with the structure of legislation and regulations. Legislation is typically organized as general rules subject to exceptions. Arguments made by applying legal rules are *defeasible*. Their conclusions can be defeated with better counterarguments. Various legal rules may conflict with each other. These conflicts are resolved using legal principals about priority relationships between rules, such as the principal of *lex superior*, which gives rules from a higher authority, such as federal law, priority over rules from a lower authority, such as state law. These properties of legal rules are well known in AI and Law and have been studied extensively. References are omitted for lack of space.

Thus we model legal rules using a defeasible rule language which has been developed especially for this purpose, as part of the Legal Knowledge Interchange Format (LKIF), and use description logic (OWL more specifically) for more limited purposes: 1) to declare the language of unary and binary predicate symbols (called concepts and roles in DL, classes and properties in OWL) of the application domain; and 2) to make assertions about these predicates, using DL axioms, which are judged to be universally true and beyond dispute in the domain.

Here we illustrate the LKIF rule language by modeling two interpretations of the concept of a derivative work in copyright law. Since opinions differ about how to interpret copyright law in the context of open source licensing issues, for example about whether or not linking to a software library creates a derivative work, an important feature of our approach is the ability to include alternative legal theories in a single model, and to construct competing arguments from these alternative legal theories.

We begin with the general rule that the copyright owner of software may license the software using any license template he chooses.

```
<rule id="DefaultLicenseRule">
  <head>
    <s pred="&oss;mayUseLicenseTemplate">
      <v>SE</v> may be licensed using
      the <v>T</v> template
    </s>
  </head>
</rule>
```

Since LKIF is an XML schema, rules are represented in XML. This particular rule has a head (conclusion) but no body (conditions). Even though the rule has no conditions, inferences made using this rule are not necessarily or universally true, but remain defeasible. We will make use of this feature to express exceptions to this general rule below.

The rule has been assigned an identifier, DefaultLicenseRule, which may be used to formulate statements about the rule. That is, rules are reified and may be reasoned about just like other objects.

The predicate symbol of the statement (proposition) in the head of the rule is specified using the `pred` attribute. Its value can be the name of a class or property in a OWL ontology, as in this example. The `&oss;` entity reference refers to the ontology, using its URI.

Declaring predicate symbols in ontologies makes it possible to divide up the model of a complex domain theory into modules, with a separate LKIF file for each module. OWL provides a way to import the classes and properties of other OWL files, recursively. Similarly, LKIF provides a way to import both LKIF and OWL files. OWL makes it easy to manage predicate symbols across the boundaries of modules and to make sure that symbols in different modules refer to the same class or property when this is desired.

The XML syntax for rules in LKIF is rather verbose and not especially readable. Fortunately, it is easy to write programs for converting XML into more readable formats. Moreover, XML structure editors exist which use style sheets to enable authors to edit XML documents directly in a more readable form. Using this feature, the above rule can be displayed in the editor as follows:

**rule** DefaultLicenseRule  
**head** SE may be licensed using the T template

We will use this more readable format for displaying LKIF rules in the remainder of this article. Next let us formulate an exception to the general rule that any license template may be used for reciprocal licenses:

**rule** ReciprocityRule  
**head**  
 not: SE1 may be licensed using the T1 template  
**domains**  
SE1 uses SE2  
SE2 has license L  
**body**  
L is reciprocal  
SE1 is derived from SE2  
 unless exists T2 : L is an instance of template T2  
 such that T1 is compatible with T2

This reciprocity rule states that a software entity, SE1, may not be licensed using a license template, T1, if the software is derived from another software entity, SE2, licensed using a reciprocal license, L, unless L is an instance of a template license, T2, which is compatible with T1. The use of domains in this rule provides control information to make use of forward chaining in the description logic reasoner, as discussed in the introduction. Notice that the conclusion of the rule is negated and that the last condition of the rule expresses a further exception, using an “unless” operator.

These two rules illustrate two kinds of exceptions. In argumentation terms, arguments constructed using the ReciprocityRule *rebut* arguments constructed using the DefaultRule and arguments which make use of the explicit exception of the ReciprocityRule, by showing that the licenses are compatible, *undercut* the reciprocity argument.



change and updates the visualization accordingly. Moreover, a “find positions” assistant, based on our formal model of abduction in Carneades argument evaluation structures [2], can be used for computing minimal sets of statements which, if accepted, would make a given statement acceptable (in) or not acceptable (out) in the argument graph. In the example, if the aim is to be able to use the EPL template, one of the positions computed is {(not (valid FSFTheoryOfLinking))}. Finally, an “instantitate scheme” assistant is available for helping users to apply argumentation schemes, such as argument from expert opinion, to construct arguments and add them to the map. One could use this tool, e.g., to add the arguments of the Free Software Foundation and Lawrence about the validity of their respective theories of linking.

## 5. CONCLUSION

We have illustrated features of the Carneades argumentation system with a prototype legal application for analyzing software licensing issues. To our knowledge, no other argumentation or rule-based system currently provides the combination of tools required for this application: 1) automatic argument construction from a knowledge base of strict and defeasible rules; 2) argument mapping; 3) argument evaluation; 4) interactive construction of arguments using argumentation schemes; 5) exploration of effects of alternative legal theories; and 6) computation of positions, using abduction. The source code of the application is freely available, as open source software, and we offer it as a possible benchmark problem to the AI and Law community.

## 6. ACKNOWLEDGEMENTS

This work was partially funded by the European projects Qualipso (IST-FP6-034763) and IMPACT (IST-FP7-247228). An earlier version of this paper was presented at the Jurix 2010 workshop on Modeling Legal Cases and Legal Rules.

## 7. REFERENCES

1. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., and Patel-Schneider, P., eds. *The Description Logic Handbook — Theory, Implementation and Applications*. Cambridge University Press, 2003.
2. Ballnat, S. and Gordon, T.F. Goal Selection in Argumentation Processes — A Formal Model of Abduction in Argument Evaluation Structures. *Proceedings of the Third International Conference on Computational Models of Argument (COMMA)*, IOS Press (2010), 51-62.
3. Bing, J. Uncertainty, Decisions and Information Systems. In C. Ciampi, ed., *Artificial Intelligence and Legal Information Systems*. North-Holland, 1982.
4. Brewka, G., Dunne, P.E., and Woltran, S. Relating the Semantics of Abstract Dialectical Frameworks and Standard AFs. *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-2011)*, (2011), in press.
5. Brewka, G. and Gordon, T.F. Carneades and Abstract Dialectical Frameworks: A Reconstruction. *Proceedings of the Third International Conference on Computational Models of Argument (COMMA)*, IOS Press (2010), 3–12.
6. ESTRELLA Project. The Legal Knowledge Interchange Format (LKIF). 2008.
7. Fiedler, H. Expert Systems as a Tool for Drafting Legal Decisions. In A.A. Martino and F.S. Natali, eds., *Logica, Informatica, Diritto*. Consiglio Nazionale delle Ricerche, Florence, 1985, 265-274.
8. Gijzel, B.V. and Prakken, H. Relating Carneades with abstract argumentation. *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-2011)*, (2011), in press.
9. Gordon, T.F., Prakken, H., and Walton, D. The Carneades Model of Argument and Burden of Proof. *Artificial Intelligence 171*, 10-11 (2007), 875-896.
10. Gordon, T.F. and Walton, D. Proof Burdens and Standards. In I. Rahwan and G. Simari, eds., *Argumentation in Artificial Intelligence*. Springer-Verlag, Berlin, Germany, 2009, 239-260.
11. Gordon, T.F. and Walton, D. Legal Reasoning with Argumentation Schemes. *12th International Conference on Artificial Intelligence and Law (ICAIL 2009)*, ACM Press (2009), 137-146.
12. Gordon, T.F. Constructing Arguments with a Computational Model of an Argumentation Scheme for Legal Rules – Interpreting Legal Rules as Reasoning Policies. *Proceedings of the Eleventh International Conference on Artificial Intelligence and Law*, (2007), 117-121.
13. Gordon, T.F. Visualizing Carneades Argument Graphs. *Law, Probability and Risk 6*, 2007, 109-117.
14. Gordon, T.F. An Overview of the Carneades Argumentation Support System. In C.W. Tindale and C. Reed, eds., *Dialectics, Dialogue and Argumentation. An Examination of Douglas Walton’s Theories of Reasoning*. College Publications, 2010, 145-156.
15. Gordon, T.F. Report on a Prototype Decision Support System for OSS License Compatibility Issues. 2010.
16. Governatori, G. On the Relationship between Carneades and Defeasible Logic. *Proceedings of the International Conference on Artificial Intelligence and Law (ICAIL-2011)*, (2011), in press, this volume.
17. McCarty, L.T. Some Arguments About Legal Arguments. *International Conference on Artificial Intelligence and Law*, (1997), 215-224.
18. Walton, D., Reed, C., and Macagno, F. *Argumentation Schemes*. Cambridge University Press, 2008.