

Elements of Argumentation

Philippe Besnard and Anthony Hunter

**The MIT Press
Cambridge, Massachusetts
London, England**

© 2008 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

For information about special quantity discounts, please email special_sales@mitpress.mit.edu

This book was set in Times New Roman and Syntax on 3B2 by Asco Typesetters, Hong Kong.

Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Besnard, Philippe, 1958–

Elements of argumentation / Philippe Besnard and Anthony Hunter.

p. cm.

Includes bibliographical references and index.

ISBN 978-0-262-02643-7 (hardcover : alk. paper) 1. Artificial intelligence. 2. Reasoning.

3. Logic. 4. Information technology. I. Hunter, Anthony, 1962– II. Title.

Q335.B489 2008

006.3—dc22

2007032257

10 9 8 7 6 5 4 3 2 1

1 Nature of Argumentation

Argumentation normally involves identifying relevant assumptions and conclusions for a given problem being analyzed. Furthermore, it often involves identifying conflicts, resulting in the need to look for pros and cons for particular conclusions.

Argumentation is a vital form of human cognition. Constantly in our daily lives, we are confronted with information that conflicts, and we are forced to deal with the resulting inconsistencies. Often we do this subconsciously: As a mental reflex, we weigh conflicting information and select some items of information in preference to others. Some of the time, we deal with conflicting information in a more conscious way. For example, if we are making a big decision, we may have in mind some of the key arguments and counterarguments. Consider a decision on where to go for a long holiday or a decision on a house to buy. Here, there is a list of options with pros and cons for each option. And when we are not sure about inconsistencies in our information, we may try to seek better information, or to seek advice, in order to resolve the inconsistency.

Professionals routinely undertake argumentation as an integral part of their work. Consider diverse types of professional, such as clinicians, scientists, lawyers, journalists, and managers, who have to identify pros and cons for analyzing situations prior to presenting some information to an audience and/or prior to making some decision. Here, many conflicts are consciously identified in the available information, and then, depending on the task being undertaken, appropriate arguments and counterarguments are constructed.

Argumentation may also involve chains of reasoning, where conclusions are used in the assumptions for deriving further conclusions. Furthermore, the task of finding pros and cons may be decomposed recursively. Thus, counterarguments may be identified that conflict with the assumptions of an argument.

In this chapter, we provide an informal coverage of the nature of argumentation. For this, we consider some definitions for basic concepts for argumentation, for the kinds of information used in argumentation, and for the kinds of agent involved in argumentation. During the course of the chapter, we will provide motivation of some of the key elements of argumentation that we plan to formalize in the rest of this book.

1.1 Basic Concepts for Argumentation

We start by providing some simple informal definitions for argumentation. We will expand on these definitions in this chapter, and then, in subsequent chapters, we will explore formal definitions for these concepts.

Argument An argument is a set of assumptions (i.e., information from which conclusions can be drawn), together with a conclusion that can be obtained by one or more reasoning steps (i.e., steps of deduction). The assumptions used are called the **support** (or, equivalently, the **premises**) of the argument, and its conclusion (singled out from many possible ones) is called the **claim** (or, equivalently, the **consequent** or the **conclusion**) of the argument. The support of an argument provides the reason (or, equivalently, **justification**) for the claim of the argument.

Contradiction One formula contradicts another formula if and only if the first negates the second. In other words, two formulae contradict if and only if they are mutually inconsistent. For example, using classical logic, if we have a claim $\alpha \vee \beta$, then the claim $\neg\alpha \wedge \neg\beta$ negates it. Similarly, a formula α contradicts a set of formulae Γ iff $\Gamma \cup \{\alpha\}$ is inconsistent. For example, using classical logic, $\alpha \vee \beta$ contradicts Γ when Γ is $\{\neg\beta, \alpha \rightarrow \beta\}$.

Rebutting argument A rebutting argument is an argument with a claim that is the negation of the claim of another argument. In other words, if an argument states that β holds, a rebutting argument takes the position that the negation of β holds, hence rebutting the argument for β . Thus, an argument A_1 that rebuts another A_2 (so A_1 is a rebutting argument) is such that the claim of A_1 contradicts the claim of A_2 . For example, using classical logic, if A_1 has the claim α , and A_2 has the claim $\neg\alpha$, then A_1 and A_2 rebut each other.

Undercutting argument An undercutting argument is an argument with a claim that contradicts some of the assumptions of another argument. Assuming classical logic, suppose an argument has a support that in-

cludes the information that β holds, and the information that $\beta \rightarrow \alpha$ holds, and the claim that α holds, then an example of an undercutting argument would be an argument with a claim that is the negation of β (i.e., $\neg\beta$) or the negation of $\beta \rightarrow \alpha$ (i.e., $\neg(\beta \rightarrow \alpha)$).

Counterargument Given an argument A_1 , a counterargument is an argument A_2 such that either A_2 is a rebutting argument for A_1 or A_2 is an undercutting argument for A_1 .

Argumentation This is the process by which arguments and counterarguments are constructed and handled. Handling arguments may involve comparing arguments, evaluating them in some respects, and judging a constellation of arguments and counterarguments to consider whether any of them are warranted according to some principled criterion.

For argumentation, we may also assume that each argument has a **proponent**, who is the person (or group of people) putting forward the argument, and that each argument has an **audience**, who is the person (or group of people) intended as the recipient(s) of the argument. To illustrate these concepts, consider the following example.

Example 1.1.1 Consider two people, Charlie and James, working in a newspaper office. Charlie makes the following argument to James. So Charlie is the proponent of the argument and James is the audience of the argument.

Claim We can publicize that Simon Jones is having an affair.

Support Simon Jones is a public person, so we can publicize details about his private life.

In response, James makes the following counterargument to Charlie. So James is the proponent of the argument, and Charlie is the audience.

Claim Simon Jones is no longer a public person.

Support Simon Jones just resigned from the House of Commons; hence, he is no longer a public person.

In order to investigate examples such as the above, we can start with a significant development by Stephen Toulmin [Tou58]. For this, Toulmin identifies the importance of a layout for an argument. He shows that to analyze an argument, it is necessary to identify the key components of the information in terms of the roles played within the argument. These components are summarized as follows:

Facts The term “fact” is used by different authors in different ways. Here we assume a fact is an item of information that is specific to a given context. For example, consider a doctor advising a patient. Facts are information on a given patient, such as *name*, *age*, and *blood pressure*. This information is only applicable to that patient. This contrasts with knowledge, in the form of perhaps defeasible rules, that can be used on all patients, such as *If a patient has high blood pressure and is middle-aged, then prescribe a low sodium diet.*

Warrant This is the part of the argument that relates facts to qualified claims. A warrant captures a form of defeasible rule (a rule that is normally valid, when the required facts hold, but in exceptional circumstances, it may fail to hold): Essentially, it says that if the required conditions (represented by the facts) hold, then there is a reason to accept the qualified claim. For this setup, we can regard the facts plus the warrant as the support for an argument.


Backing A backing is some kind of justification for a warrant. It provides an explanation for why the warrant is a reason to accept the qualified claim. Justifications may be based on diverse criteria such as belief, law, authority, ethics, morals, or aesthetics.

Rebuttal A rebuttal captures the circumstances that would be regarded as exceptions for a warrant. In other words, it captures the reasons that would render the warrant as not holding. Thus, if the facts for the rebuttal hold, then we have a rebutting argument, and hence a counterargument, for the argument based on the warrant.

Qualified claim A qualified claim is a conclusion that can be drawn if the warrant holds and the rebuttal does not hold. In a sense, the facts plus the warrant imply the claim.

An example of an argument, conforming to Toulmin’s layout, is given in figure 1.1. Here, we see that from the fact that *Harry was born in Bermuda*, we have the qualified claim that *Harry is a British subject*, unless *both his parents were aliens* or *he has become a naturalized American* or etc.

In a sense, Toulmin’s approach to argumentation in terms of the layout of arguments is analogous to using a classification system or decision tree: Given some facts, we can decide whether the qualified claim holds by checking whether the facts and warrant hold and the rebuttal does not hold. The approach is structural and, in a sense, logical. However, it lacks mechanisms for constructing and manipulating the graphs. It is,

**Figure 1.1**

An example of argumentation taken from [Tou58].

in a sense, static or fixed. Furthermore, it is text-based, and so it requires some interpretation to use it. This makes it difficult to automate reasoning with it.

To illustrate shortcomings of Toulmin's proposal if we want to use automated reasoning, suppose we have some individual *Mr. Jones*, and we want to determine whether *Mr. Jones is a British subject*; then we need to determine whether *Mr. Jones was born in Bermuda* and whether *both his parents were aliens* or *he has become a naturalized American* or etc. This is not undertaken in any formal language; it is done in natural language and therefore is subject to the problems of ambiguity that arise with the usage of natural language in computing.

For example, suppose we have the fact that *Mr. Jones was born in France*; in order to automate the reasoning that would prohibit the usage of this particular warrant for *Mr. Jones*, we would need to state explicitly that *France is not part of Bermuda*. Of course, this is trivial for a human, but for computing it may require much commonsense knowledge being explicitly specified, and we would need some knowledge representation and reasoning formalism for capturing and automating the use of this commonsense knowledge. As another example, now suppose instead we have the information that *Mr. Jones was born in Hamilton*; again, we need commonsense knowledge and automated reasoning that would allow the usage of the warrant for the individual *Mr. Jones*.

Thus, Toulmin's layout of arguments gives us some important concepts, including warrant, backing, rebuttal, and qualified claim, for describing arguments. We can see these concepts featuring in numerous examples. However, the approach does not just provide a comprehensive account of the logic of argumentation, and furthermore, the approach does not address many important questions on how to automate the construction or use of layouts of arguments. Nonetheless, it would be reasonable to suggest that Toulmin's layout of arguments is an antecedent to many formal approaches to argumentation in artificial intelligence, though many of the formal proposals deviate significantly from this starting point.

1.2 Information Involved in Argumentation

At the core of argumentation is the need for information. If we have no information, we have no arguments, except perhaps tautologies.

Potentially, argumentation can be based on any kind of information. In the following, we consider some informal delineations of types of information. Our aim in presenting these classifications of information is solely to support our presentation of argumentation. We do not wish to suggest that these classifications are contributions in their own right.

Information can be described as being either certain or uncertain as delineated in the following descriptions:

Certain (or categorical) information This is information that is treated as absolutely correct. It is straightforward to treat as certain a mathematical definition or commonly-known knowledge like *the capital of France is Paris*. However, a broader range of examples includes information where the possibility of doubt is so small that the information can be regarded as certain, like *tomorrow the sun will rise*.

Uncertain information This is information that is not certain. Most information is uncertain to some degree. Deciding on whether information is certain or uncertain often depends on the circumstances in which the information is used and evaluated—for example, *the length of the banana is 15 cm*, *it will rain in New York tomorrow*, and *Mr. Jones has had a mild infarction*.

Deciding whether information is certain or uncertain can depend on the application. For example, it is reasonable to assume that *every day the sun will rise* is certain, but there is a minuscule chance of a cosmic catastrophe that would result in the sun's failing to rise tomorrow. Thus, for example, if we consider going for a walk in town tomorrow, it is reasonable to assume that *the sun will rise tomorrow* (though it may be obscured by clouds), but if we consider the future of the universe, we may consider that it is not certain that the sun will rise on any particular day in the future.

We do not need to get sidetracked here into the nature of uncertainty. Rather, we suggest that the reader may wish to adapt the definitions for these terms according to their needs. For more information on the nature of uncertainty see [KC93, PH98].

Information (both certain and uncertain) can also be described as being one of objective, subjective, or hypothetical, as follows:

Objective information This is information that comes from a “reliable source” or can be observed, measured, or verified by everyone involved in the argumentation. For example, consider *a clinical trial for a new drug treatment where 90% of the group of patients with the new treatment survive after five years and 20% of the group of patients with the control treatment survive after five years*. This is objective information. However, just because information is objective, this does not mean that it is necessarily correct or consistent. Errors, and so forth, can occur in obtaining objective information. A possibility for a major error could be that the way the patients were selected for the trial might unknowingly allow for a selection that would respond very well to the treatment, whereas in a wider population of patients with the disease, the success rate may be substantially lower. Thus, if the above value of 90% was used for the population in general, it would be erroneous. In general, inconsistency in information has a diverse etiology, but in objective information it often arises from errors. Consider, for example, two honest witnesses to a bank robbery: Overall, they may give useful information on the event, but they may give quite conflicting descriptions of the getaway car.

Subjective information This is information that comes as beliefs or opinions from some of those involved in the argumentation. This is not necessarily consistent information. An example of subjective information may arise when an oncologist is advising a patient on options for treatment plans. Here, the oncologist may present arguments for and against some options involving combinations of radiotherapy, chemotherapy, and surgery. Much of the information used by the oncologist in the supports for the arguments will be objective, for example, scientific knowledge about the relative efficacies of treatments, but it may also involve preferences expressed by the patient about different drugs or about the weight the patient would put on quality of life over overall survival chances.

Hypothetical information This is information that is assumed for the sake of constructing arguments of interest. It is not necessary for hypothetical information to be true. It may even be such that it is unlikely to be true now or in the future. It may still be useful to consider it as part of the assumptions for argumentation if one wants to explore possibilities. Thus, some hypothetical information may be described as speculative information. For example, it is unlikely that the sea level will rise by 50 cm in the next twenty years, but it would be useful for a government to consider the possibility in their coastal areas in order to consider whether or not they are adequately prepared for flooding. This could be done by assuming the hypothetical information that *the sea level will rise by 50 cm in 20 years*, and then arguments could be constructed for and against the conclusion that they are adequately prepared. As another example, the government may wish to extend this civil emergency planning to consider the hypothetical information that *a UFO will land in the country next week*. Hypothetical information may even include information that the proponent regards as false (i.e., the proponent regards it as fallacious information). Consider, for example, how a sophist may construct an argument for a claim of interest.

Deciding whether information is objective, subjective, or hypothetical can also depend on the application. Again, we do not need to get side-tracked here into more precise definitions for these categories. We present these categories only to indicate the range of situations for which we may wish to formalize argumentation. There are also other dimensions that we could consider for describing information (including epistemics, deontics, vagueness, and particular kinds of uncertainty such as probability and possibility) based on developments in the knowledge representation and reasoning literature.

1.3 Agents Involved in Argumentation

Central to conceptualizing argumentation is that argumentation involves agents and groups of agents. This is for considering both the proponent and the audience of each argument.

First, we will delineate a notion of agent and entity, and then we will define the notions of monological argumentation and dialogical argumentation in terms of agents and entities involved. To support this goal, we will adopt the following informal definitions (for a comprehensive and formal conceptualization of agents and associated notions, see [Woo01]):

Agent An agent is an autonomous, proactive, and intelligent system that has some role. Examples of kinds of agent include lawyers, clinicians, and journalists. Further examples of types of agent include voters in an election, readers of a newspaper, jury members, and patients in a hospital. We may wish to also think of some software systems as agents if they display sufficiently significant intelligence, autonomy, and proactiveness.

Entity An entity is composed of a set of agents that in concert have some role. A simple example of an entity is a board of directors for a company, where each agent in the entity is a director. The agents in an entity may be heterogeneous. In other words, different agents in an entity may have different roles. For example, a court is an entity that is composed of a judge, a prosecution lawyer, a defense lawyer, witnesses, a defendant, and jury members. These are agents with different roles, and in concert they have the role of conducting a trial of the defendant. Another example of an entity is an audience for a political speech. Here the audience may be composed of agents who each have a political standpoint, and so in this case the role of the entity is only to be an audience to the political speech. A third example of an entity is a group of scientists, working on a research project, who publish a scientific paper.

Thinking of argumentation in terms of agents allows us to formalize the different roles that agents can play in different kinds of argumentation. In order to further delineate our concerns in this book, we need to briefly describe the monological and the dialogical views on argumentation:

Monological A single agent or entity has collated the knowledge to construct arguments for and against a particular conclusion. This involves collating both categorical and uncertain information. Furthermore, this

may include objective information (e.g., externally measured or verifiable information, information obtained from reliable third-party sources, etc.), subjective information (e.g., beliefs, aesthetics, etc.), and hypothetical information. The knowledge may come from heterogeneous sources. After constructing the arguments, the entity may then draw some conclusion on the basis of the assembled arguments. The emphasis of the monological view is on how to construct the arguments and how to draw conclusions from the assembled arguments. Monological argumentation can be viewed as an internal process for an agent or an entity with perhaps a tangible output (e.g., an article or a speech or a decision). In monological argumentation, there is no representation of the dialogue between the agents or entities involved. However, the knowledge used to construct the support for one or more arguments may have been obtained from a dialogue.

Dialogical A set of entities or agents interact to construct arguments for and against a particular claim. If an agent offers an argument, one or more of the other agents may dispute the argument. Agents may use strategies to persuade the other agents to draw some conclusion on the basis of the assembled arguments. The emphasis of the dialogical view is on the nature of the interactions and on the process of building up the set of arguments until the agents collectively reach a conclusion. Dialogical argumentation can be viewed as incorporating monological argumentation, but in addition, dialogical argumentation involves representing and managing the locutions exchanged between the agents/entities involved in the argumentation.

In a sense, monological argumentation is a static form of argumentation. It captures the net result of collating and analyzing some conflicting information. In contrast, dialogical argumentation is a dynamic form of argumentation that captures the intermediate stages of exchanges in the dialogue(s) between the agents and/or entities involved.

Nonetheless, monological and dialogical argumentation involve a proponent and an audience. Some agent or entity provides each argument, and some agent or entity is the intended audience for that argument—though, of course, the proponent and audience for an argument may be the same agent or entity, particularly in the case of monological argumentation.

To illustrate the difference between monological and dialogical argumentation, we consider some examples. For monological argumentation, we list some situations for static argumentation and the kinds of agent or entity that are responsible for producing that argumentation:

- A newspaper article by a journalist.
- A political speech by a politician.
- A political manifesto by a political party.
- A review article by a scientist.

For dialogical argumentation, we list some situations for dynamic argumentation and the kinds of agent or entity that are responsible for that argumentation:

- Lawyers arguing in a court.
- Traders negotiating in a marketplace.
- Politicians debating about new legislation.
- Governments negotiating a new world trade agreement.
- Family members arguing over who should do the washing up.

Ultimately, both monological and dialogical argumentation aim for some final result from the process, but in monological argumentation, the emphasis is on the final result, whereas in dialogical argumentation, the emphasis is the process as represented in terms of dialogue exchanges. This obviously has important ramifications for formalizing argumentation. To formalize dialogical argumentation, a lot of extra machinery is required to model or automate the role of the agents involved.

It is clear that there are a variety of roles for monological argumentation depending on the kind of information used and on the aim of the presenter of the arguments. The following breakdown is only meant to indicate the diversity of roles for monological argumentation; it is not meant to be a “definitive classification.”

Factual argumentation Use just objective information with the aim of informing the audience about some verifiable information—for example, a scientific review. Here, we assume that there is no hidden bias in how the argumentation is undertaken.

Positional argumentation Use objective information, subjective information, and hypothetical information with the aim of informing the audience of the presenter’s beliefs—for example, a newspaper opinion article.

Persuasive argumentation Use objective information, subjective information, and hypothetical information (including possibly fallacious information) with the aim of persuading the audience to do something—for example, a political speech, a team pep talk, or a sales pitch.

Provocational argumentation Use objective information, subjective information, and hypothetical information (including possibly fallacious information) with the aim of provoking the audience of some hypothetical situations for entertainment, to invoke further thinking, or to map extremes in a space—for example, a newspaper opinion article, a think-tank pamphlet, or an academic article. Provocational argumentation can also be used as entertainment, such as in satire and in sophism.

Speculational argumentation Use objective information, subjective information, and hypothetical information (including speculative information) with the aim of informing the audience about a possible scenario for explaining some past event or some possible future event—for example, a risk management scenario or an academic article.

Monological argumentation has directionality. In other words, when an agent constructs some arguments and counterarguments, there is normally an intended recipient in some sense. The intended audience can range from one particular agent through to the global audience (i.e., anybody).

We regard all monological argumentation as either argumentation by the proponent for the proponent (auto-argumentation), and so the proponent and the intended audience are the same, or argumentation by the proponent for one or more other agents (one-to-many argumentation):

Auto-argumentation This is argumentation for agents(s) to identify key arguments and counterarguments for their own use, such as for problem analysis prior to making a decision. For example, for most of us, when we buy a house, we have a limited budget, a list of features we would like, and a list of features we would dislike. It is often the case that we can narrow the choice down to a few possible houses that are available, and none of them are perfect. Each may lack some of the features we would like, and each may have some features we would dislike. In other words, each of the houses on the shortlist is inconsistent with our requirements. However, because we have to make a choice, we can consider the pros and cons for each possible house. Auto-argumentation could also be called self-argumentation.

One-to-many argumentation This is argumentation by an agent or entity for distribution to other agents or entities—for example, a newspaper article by a journalist, a lecture by a scientist, or a speech by a politician. Of course, one-to-many argumentation does not have to involve professionals. Consider, for example, a child making a case to his or her parents

for a higher allowance. One-to-one argumentation (a proponent presenting an argument to an audience of exactly one agent) is a special case of one-to-many argumentation.

Even though monological argumentation is argumentation that involves just one proponent, which may be an agent or an entity, it may summarize information that has come from debates or discussions, as well as other sources of information.

To illustrate this, consider a situation where a clinician and a patient discuss options for an oncology treatment plan, taking into account relevant medical knowledge together with the patient's preferences and personal circumstances. The meeting involving these two agents would involve dialogical argumentation. However, when they have exchanged the necessary information, and perhaps debated some of the major issues of concern, they could then bring together the major points for the patient, highlighting them in the form of the pros and cons of the key options, prior to the patient's making necessary decisions and/or giving the necessary authorization. This representation of the key information gleaned during the meeting in the form of pros and cons is a form of monological argumentation. Furthermore, since the doctor and the patient come together to act as an entity in collating the arguments and counterarguments, it is a form of auto-argumentation.

We will see further examples, in this book, of how we can represent the key information gleaned during a discussion or debate using monological argumentation. We may choose to think of there being a "third agent" who is collecting information from the dialogical argumentation and using it to undertake monological argumentation, without necessarily attributing the gleaned information to any of the sources and without representing any of the history of the dialogue.

1.4 Requirements for Formalizing Argumentation

The overall aim of this book is to present formalizations of aspects of monological argumentation. In so doing, we will formalize key elements of practical argumentation. By practical argumentation, we mean argumentation that reflects more closely argumentation as practiced by agents in the real world.

In our coverage, we will consider how abstract argumentation and logical argumentation provide important foundations for formalizing monological argumentation. We will also see shortcomings in the basic versions

of these proposals for capturing practical argumentation. In order to consider these proposals systematically, we now sketch the requirements we have for formalizing practical monological argumentation in this book. As part of presenting solutions to these requirements, during the course of the book, we will be conceptualizing some of the key elements of argumentation:

Presentation of arguments We want to be able to present an exhaustive display of the constellation of arguments and counterarguments relevant to a particular claim. This should act as an inventory of all the different ways that the conclusion can be inferred from the assumptions and all the different ways that counterarguments can be inferred from the assumptions. Given a particular claim and a “knowledgebase” (from which we find the supports for arguments and counterarguments), we want to be able to automate the construction of each constellation.

Analysis of intrinsic factors Given a constellation of arguments and counterarguments relevant to a particular claim, we want to analyze the nature and type of conflicts that arise in the constellation. We also want to be able to annotate the constellation with information about the results of analyzing intrinsic factors.

Analysis of extrinsic factors For a constellation of arguments and counterarguments relevant to a particular claim, we want to analyze the quality of the arguments and counterarguments from the perspective of a representative (or stereotypical) member of the audience. This includes considering how believable the arguments are for the representative and what the impact is for the representative. We also want to be able to annotate the constellation with information about the results of analyzing extrinsic factors.

Selection of arguments Given the ability to undertake analyses of intrinsic and extrinsic factors, we want principled techniques for selectivity in the choice of arguments and counterarguments used in a constellation. The net result is that, as an alternative to an exhaustive presentation of arguments and counterarguments, we obtain a more focused constellation of arguments and counterarguments tailored for the intended audience. Being selective means that the argumentation can be made more believable and have higher impact for the intended audience.

Judgment of constellations For a constellation of arguments and counterarguments relevant to a particular claim, we want principled criteria for suggesting whether the claim is warranted or unwarranted.

Reformation of constellations Given a constellation of arguments and counterarguments relevant to a particular claim, we want principled means for reforming (i.e., restructuring) the arguments by, for example, merging arguments with logically equivalent supports.

We want to be able to present an exhaustive display of arguments and counterarguments relevant to the conclusion as output because we want the user to decide what to do with the information. We do not want to develop a black box for outputting conclusions; rather, we want to output the key arguments and highlight the key conflicts. If we consider some of the kinds of monological argumentation that we are interested in capturing, such as newspaper articles, political speeches, and scientific research papers, it is clear that the information assumed, and the way it is put together, is as important as, if not more important than, the conclusion obtained.

However, there is also normally the need for arguments to be apposite for the intended audience. Consider an article in a current affairs magazine: Only a small subset of all possible arguments that the journalist could construct from his or her own knowledgebase is used. The journalist regards some arguments as having higher impact or as being more believable for the intended audience or more relevant than others and so makes a selection. This need for appositeness is reflected in law, medicine, science, politics, advertising, management, and just ordinary everyday life.

Thus, taking the audience into account means that there has to be some selectivity of the arguments presented to them. Numerous formal theories of argumentation exercise selectivity on grounds of certainty and preference as viewed from the presenter's perspective, but the audience's perspective is largely ignored. We want to formalize these in terms of knowledge about the audience. We will argue that being selective in argumentation improves the constellation of arguments and counterarguments by making it more interesting and more believable.

We can think of the requirements giving the user a range of options. The user of an argumentation system can choose to have an exhaustive display of a constellation of arguments and counterarguments, or the user can choose to have a selective display of a constellation of arguments and counterarguments, based on a particular audience. In either case, the user can choose to annotate the constellation with information coming from analyzing intrinsic and extrinsic factors arising in the constellation.

1.5 Frameworks for Formalizing Argumentation

If we want to handle arguments systematically, then we need a “formalization” of argumentation. Many professions implicitly or explicitly explore these issues and, indeed, put the systematic use of arguments at the heart of their work. Consider, for example, the legal, medical, and journalistic professions.

However, in this book, we want to go beyond the systematic handling of arguments: We want to handle arguments automatically, and we want the techniques to scale up to handling substantial and complex problems. This calls for more detailed formalizations with algorithms. Furthermore, if we want predictable behavior for our argumentation systems, then we need theoretical properties and empirical results. This, in turn, will call for a sophisticated and precise understanding of the principles of argumentation, which, in turn, calls for richer and deeper theoretical formalisms.

Classical logic is appealing as a starting point for argumentation: The representation is rich and the reasoning powerful. Furthermore, it can be argued that classical reasoning captures some of the important ways that people undertake logical reasoning: For example, modus ponens, modus tollens, and disjunctive syllogism. However, the appeal of classical logic extends beyond the naturalness of representation and reasoning. It has some very important and useful properties that mean that it is well-understood and well-behaved and that it is amenable to automated reasoning.

In classical logic, statements are represented by formulae. Both assumptions and conclusions are represented by formulae, and the language for assumptions and conclusions is the same. Let Δ be a set of formulae, let $\vdash$ be the classical consequence relation, and let α be a formula; then $\Delta \vdash \alpha$ denotes that α is an inference (i.e., a conclusion) from Δ using classical logic. In Appendix C, we provide a review of the syntax and semantics for classical logic.

However, there is a key concern if we are to use classical logic for argumentation. We have already acknowledged that argumentation involves considering conflicting (i.e., inconsistent) information. If the knowledge we have for constructing arguments is consistent, then we will not be able to construct conflicting arguments, and hence we will not have recourse to argumentation. Unfortunately, inconsistency causes problems in reasoning with classical logic.

In classical logic, any inference can follow from an inconsistent set of assumptions. A useful definition of inconsistency for a set of assumptions Δ is that if $\Delta \vdash \alpha$ and $\Delta \vdash \neg\alpha$, then Δ is inconsistent. A property of classical logic is that if Δ is inconsistent, then for any β in the language, $\Delta \vdash \beta$. This property results from the following proof rule, called *ex falso quodlibet*, being a valid proof rule of classical logic.

$$\frac{\alpha \quad \neg\alpha}{\beta}$$

Thus, inconsistency causes classical logic to collapse. No useful reasoning follows from an inconsistent set of premises. It can be described as exploding, or trivialized, in the sense that all formulae of the language are consequences of an inconsistent set of assumptions. From a semantic perspective, there are no models of a set of inconsistent formulae.

Partly in response to the issue of inconsistency arising in argumentation, there have been three main approaches to formalizations for argumentation, namely, abstract systems, defeasible systems, and coherence systems. The first two approaches use formalisms that are, in key respects, much less expressive (in terms of the complexity of information that can be represented and in the complexity of the inferences that can be drawn) when compared with classical logic, thereby circumventing the problem of inconsistency as manifested by *ex falso quodlibet*, and the third approach adopts a simple strategy to ameliorate the problem of inconsistency. We delineate these three approaches as follows:

Abstract systems These are based on the seminal proposal by Dung [Dun95] that assumes a constellation of arguments and counterarguments can be captured by a set of arguments and a binary “attacks” relation between pairs of arguments. The attacks relation captures the situation where one argument undermines the credibility of another. This setup can be viewed as a graph, with each node representing an argument and each arc representing an “attacks” relationship. Thus, the constellation, represented by the graph, is the starting point. It is not constructed from a knowledgebase. Reasoning with the graph is based on finding coalitions of arguments such as a coalition of arguments that do not attack each other and that attack any argument that attacks any member of the coalition. We review abstract systems in chapter 2.

Defeasible systems There are a number of proposals for defeasible logics. The common feature for these logics is the incorporation of a

defeasible implication into the language. Defeasible logics have their origins in philosophy and were originally developed for reasoning problems similar to those addressed by nonmonotonic logics in artificial intelligence. Arguments can then be defined as chains of reasons leading to a conclusion with consideration of potential counterarguments at each step. With the explicit structure in the chains of reasoning, diverse notions of defeat can be conceptualized. Once nonclassical notions of implication are introduced into the language, giving rise to either subclassical systems (i.e., systems weaker than classical logic) or superclassical systems (i.e., systems stronger than classical logic), an interesting range of issues arise for refining the notion of an argument, a counterargument, an undercut, a rebut, and so on. A number of these defeasible systems construct arguments logically and then evaluate sets of them as an abstract system (each logical argument is a node in the graph, and if an argument rebuts or undercuts another, then this is represented by an arc in the graph). In this way, a defeasible system can “instantiate” an abstract system, or equivalently, the abstract system provides a “semantics” for the defeasible system. We review some of these issues in chapter 8.

Coherence systems One of the most obvious strategies for handling inconsistency in a knowledgebase is to reason with coherent (i.e., consistent) subsets of the knowledgebase. This is closely related to the approach of removing information from the knowledgebase that is causing an inconsistency. In coherence systems, an argument is based on a consistent subset of an inconsistent set of formulae—the inconsistency arises from the conflicting views being represented. Further constraints, such as minimality or skeptical reasoning, can be imposed on the consistent subset for it to be the support for an argument. The most common choice of underlying logic for coherence systems is classical logic, though other logics such as modal, temporal, spatial, or description logics are possible. While coherence systems, based on classical logic, give substantial expressibility in order to capture a wide range of monological argumentation situations, there may be computational and applicational advantages of using argumentation systems based on simpler defeasible logics.

We will cover abstract systems in detail in chapter 2, and we will cover defeasible systems in detail in chapter 8. The main framework that we will present in this book can be regarded as a coherence system. We will introduce this in chapter 3 and develop it during subsequent chapters. In this way, we will use the logic-based approach presented in chapter 3 as the vehicle to isolate key elements of argumentation during the subse-

quent chapters, and we believe that many of these elements can be seen in all logic-based approaches to argumentation.

1.6 Discussion

In this chapter, we have attempted to delineate some of the basic concepts that are part of argumentation. In addition, we have considered our focus in this book on formalizing monological argumentation. This type of argumentation is central to other forms of argumentation. It is an interesting subject of study in its own right, and it offers much potential as part of technological solutions for decision support, multi-agent systems, and computational linguistics.

For the logic-based frameworks we present in this book, we are assuming that the input for a system based on monological argumentation is a knowledgebase, together with a claim of interest, and the output is a constellation of arguments and counterarguments. This constellation may have been subject to selectivity according to analysis based on intrinsic and/or extrinsic factors, and it may have been annotated with information on some of these analyses. Via the presentation of this framework, we aim to delineate some of the key elements of argumentation.

In section 1.4, we sketched some of the key requirements for formalizing practical argumentation, by which we mean we have delineated some of the key features of argumentation undertaken by real-world agents with the aim of being able to capture these features in a logic-based setting. In order to investigate the issues surrounding this aim, we have decoupled practical argumentation from the wider aims of an agent, such as planning and acting in the real world. Though, as we suggest later, and as other authors have recognized (e.g., [AP05a, AP05b]), we will need to consider these wider issues to more fully understand practical argumentation.

1.7 Bibliographic Notes

Numerous textbooks have explained how individual arguments, as found in the real world, can be represented and analyzed by classical logic. The focus of these books is on what constitutes a valid logical argument. A paragon of such a textbook by Fisher [Fis88] includes many examples of arguments originally presented in philosophy and in politics using free text, together with a comprehensive explanation of how they can be translated into propositional logic. However, these textbooks tend

to circumvent the more difficult issues of inconsistency, conflict, and counterarguments.

To address these more difficult issues, an excellent starting point for considering monological argumentation is Toulmin's book [Tou58], with developments of Toulmin's approach having been reviewed by van Eemeren et al. [vGK87]. Toulmin's work is regarded as a precursor to much of the formal developments in argumentation systems found in the artificial intelligence field. For reviews of formalisms for argumentation systems in artificial intelligence, see [PV02, CML00].

While much interesting progress has been made in recent years on formalizing argumentation, it is clear that there are many more difficult features of argumentation that remain to be captured in some way. Within the philosophy community, there has been a line of research called "informal logic" that has been studying the nature of argumentation, and this has resulted in a number of valuable conceptualizations for the development of argumentation in artificial intelligence (see, e.g., [Wal89]).

More generally within philosophy, there has long been an interest in argumentation (for a review from a historical and sociological perspective, see [Bil87]). The initiation of the study of formal deductive reasoning is often attributed to Aristotle. He can also be said to have initiated the study of rhetoric, which in part considers how a proponent should take the audience into account when presenting arguments. While we do not wish to get sidetracked into a review of the study of rhetoric, more recently, Perelman [Per82] reworked the philosophy of rhetoric in a way that offers some valuable insights into the importance and nature of audiences for argumentation. We view Perelman's work as providing valuable background and motivation for some of the formal developments we present in this book for taking the audience into account. Still more recent background for taking the audience into account is presented in Cockcroft and Cockcroft [CC92] and Hollihan and Baaske [HB05].