

Mapping The Structure of Debate¹

Abstract: Although debate is a richly structured and prevalent form of discourse, it has received little scholarly attention. Logicians have focused on the structure of individual arguments—how they divide into premises and conclusions, which in turn divide into various constituents. In contrast, I focus on the structure of *sets* of arguments, showing how arguments are themselves constituents in high-level dialectical structures. I represent debates and positions by graphs whose vertices correspond to arguments and whose edges correspond to two inter-argument relations: “dispute” and “support,” respectively. On this basis I develop a theory of the structure of debate.

Keywords: debate, argumentation analysis, argument diagrams, argument visualization, dialectics, discourse analysis

Introduction

From 1992-1997, I took part in a project mapping out the artificial intelligence debate (Horn, Yoshimi, Deering, and McBride, 1988). It was a daunting effort. We read over 400 academic articles and books on the topic of machine intelligence, and in each case distilled whatever claims we could from the heavy prose. When the dust settled, we had extracted 805 arguments from over 380 authors and had collected these arguments into 77 “regions” of debate: e.g. “Can Computers be Creative?”, “Can Computers be Conscious?”, “Can Computers Have Free Will?” etc.

In order to organize these arguments in an intelligible way, we were forced to reckon with high level structures that have received little attention in the literature. This led us to formulate a variety of inter-argument relations (in particular, dispute and support) and to use these to describe such structures as “threads,” “debates,” “positions,”

¹ Penultimate version of a paper which appeared in *Informal Logic*, 24:1 (2004).

and “fronts.” What emerged from these efforts was a theory of the structure of debate. In what follows I outline and expand on that theory, positioning it relative to existing work in the field and suggesting directions for future research.

Debate level structures are worth studying for the obvious reason that they are pervasive—any time two or more parties trade off arguments, a debate is underway.¹ But there are other reasons. Debate-maps have heuristic utility, providing compact overviews of densely-contested realms of discourse. Not only do such overviews save time, they prevent redundant argumentation, facilitate the teaching of critical thinking skills (Van Gelder, 2003), and provide a framework for managing the kinds of cognitive complexity that have been shown to emerge in perspectively diverse environments (Antonio et al., 2004). Some have suggested that appropriate representation of ongoing debate could accelerate scientific progress (Eagleman & Holcombe, 2003).

After reviewing related work in the field, I consider the concepts of “argument” and “dispute” from the perspective of debate-structure. Then, by treating arguments as vertices and disputes as edges, I use graph theory to define threads (sets of arguments disputing one another), debates (sets of threads with a shared “root argument”), positions (sets of arguments related by another inter-argument relation: support) and fronts (where debates meet positions). Along the way I suggest ways of perspicuously depicting debates in two-dimensional spaces (e.g. chalkboards and computer screens).

It is important that a pernicious equivocation be flagged at the outset. The oft-noted ambiguity between “arguments” in the sense of premises and conclusions and “arguments” in the sense of people debating one another is especially confusing in this context, where arguments in both senses play a prominent role. I will use “argument” in

the former sense, to mean a set of statements at least one of which is a conclusion, and reserve “debate” for clashes between differing perspectives.

Background

The study of debate has its origins in classical antiquity. Dialectic was a central object of study in Greek rhetoric, Roman law, and Medieval disputation (whose vestigial remnant is the modern-day doctoral examination).² In the twentieth century such trends as "dialectics," "the new rhetoric," and “argumentation analysis” have theorized issues similar to those presented here.³ Some, especially in the last five years, have also studied the graphical representation of debate.⁴

In this section I position debate-mapping relative to these studies. This not only sharpens the scope and purpose of this study, it also suggests ways in which these different and complimentary (but easily confused) enterprises can be brought together. Debate mapping, we shall see, treats debates themselves (rather than arguments) as objects of study, with a focus on the "deep structure" which emerges when the redundancies and infelicities of spoken debate are trimmed away.

The Hierarchy of Logic. In positioning the study of debate relative to other work in the field, it is useful to view logic as a discipline which, like biology, studies hierarchically organized structures. Arguments, like organisms, can be broken down into increasingly simple parts and aggregated into increasingly complex systems. But, while biology studies its entire hierarchy from molecule to social group, logic has focused largely on the lower and middle levels of its hierarchy—arguments and their

constituents. I focus on the upper levels—providing a kind of ecology for logic—by considering how arguments aggregate into threads, debates, and positions. Note that the debate-level structures I study occupy an even higher level than what Freeman and others call “macro-structure.”⁵

It is important not to confuse argument-level structures with debate-level structures. For example, argument diagrams—a standard tool in introductory critical reasoning courses—employ graph theory similar to that employed here.⁶ But whereas argument diagrams relate premises and conclusions *within* an argument (allowing one to distinguish divergent, convergent, linked, and serial arguments, among others)⁷, debate maps relate whole arguments (allowing one to distinguish different forms of thread, debate, and position). Thus, every node on a debate-map can be represented by its own argument map, resulting in a graph of graphs.

Argumentation Analysis vs. Debate Analysis. Most studies of “argumentation analysis” which take account of debate or dialectical exchange do so from the perspective of individual arguments (Freeman, 1991; Wigmore as described by Goodwin, 2001; Toulmin, 2003; most involved in the emerging field of “Computer Supported Argument Visualization;” see Kirschner, Buckingham Shum, & Carr, 2003). The object is typically to build a case around a single claim—to manage the gathering of evidence in support of it and to show how it fares in the face of various kinds of criticism. This has obvious merit, especially to the person building the case. Moreover, the process of strengthening an argument so that it can handle disputes can lead (by a “dialectical process” reminiscent of Hegel) to more complex, subtle arguments.

While argumentation analysis focuses on a single argument and draws in a local penumbra of relevant debate, debate maps treat the debate itself as an object of study. This difference in emphasis has substantive consequences. For example, absent the debate-level vantage one is not likely to consider the “fronts” which develop between competing families of arguments, nor is one likely to inquire into the structure of “threads.” Moreover, while argumentation analyses treat single statements as basic elements, debate maps treat whole arguments as basic.

Surface structure vs. deep structure. Some theorists (Eemeren & Grootendorst, 1992; Freeley & Steinberg, 2000; Rescher, 1977) do focus on debates rather than individual arguments, and they do so in terms similar to those outlined in what follows. Rescher, for example, describes an "exfoliating tree" of disputation which "presents a sequential unfolding of point and counterpoint" (p. 5). These studies put debate structure to work in a variety of contexts: classifying fallacies more systematically than is otherwise possible (van Eemeren & Grootendort), formulating rules for critical discussion (van Eemeren & Grootendort), classifying patterns of critical exchange (Rescher), facilitating the analysis and judgment of formal debate (e.g., Freeley & Steinberg), and developing an epistemological system in which those beliefs are warranted which withstand critical scrutiny (Rescher).

Such studies focus on what I shall call the "surface structure" of debate, the process of back-and-forth which unfolds between interlocutors who argue out-loud, on-line, or in general without the time to carefully consider the conceptual structure of their exchange. Surface structure is legitimate and important as a topic of study, and debate maps can capture it, provided dispute and support are suitably redefined.⁸ However, this

study is focused on "deep structure"—that network of arguments and inter-argument relations which genuinely obtain in some contested realm of discourse.⁹

Deep structure is surface structure pruned of redundancy and irrelevancy (discussed below). Whereas surface structure has a "fast dynamic" of quickly traded point and counterpoint, deep structure has a slower dynamic—new arguments are added to a debate map only after careful analysis reveals that they are unique and genuinely relevant to their targets. The model for this form of analysis is published debate, where it is assumed that peer review and publication standards have done the work of pruning redundant and irrelevant arguments.¹⁰ Robert Horn (Horn, 2003) calls this "relatively stable argumentation," and notes that debate maps of this kind are simultaneously historical documents, insofar as they track the history of a public debate.

Some specific differences between the two kinds of analysis are as follows. Redundant arguments and irrelevant counter-arguments are common and effective in live argumentation, but are pruned from debate maps. Concessions and retractions are essential moves in an honest critical exchange, but since they do not correspond to new arguments, they need not be included in a map of deep structure.¹¹ Tentative or provisoed assertions, such as "well, I could be wrong, but it seems that your claim is problematic in the following way," are a natural part of live debate (Rescher gives a detailed analysis of these forms of move), but are not necessary on a debate map, which focuses on claims that someone is willing to put in print on the basis of careful analysis.

Deliberation vs. Debate. Some theorists (Conklin, 2003; Van Gelder, 2003) endeavor to achieve what I call "deep structure"¹² using graphs structurally isomorphic to the ones I introduce. The emphasis in these cases is *deliberation*, "a process aimed at

deciding whether some claim ought to be believed by considering relevant arguments” (Van Gelder, p. 98). Software packages are used to facilitate such deliberation.

The differences between debate mapping and computer-facilitated deliberation are largely differences of emphasis. In fact, deliberative systems can be seen as implementing the debate-maps I describe, and conversely, this paper can be seen as articulating the philosophical underpinnings of deliberative systems. However, there are differences between the two enterprises. The main difference is that deliberation tools facilitate the pursuit of rational consensus (on a policy issue, for example), generally on a scale of weeks or months, whereas debate maps describe longer-term controversies between parties who *genuinely disagree* (that is, who have, even after careful reflection, failed to achieve consensus with one another). Parties in a facilitated deliberation are not strongly committed to one or another argument—they actively consider all of them. With debate maps, by contrast, arguments are associated with committed authors. Moreover, to some extent deliberative systems fall into the “argumentation analysis” camp, insofar as they are more focused on individual arguments (with an eye towards resolution) than on debate itself. But again, these are differences of emphasis: debate maps do intend to facilitate deliberation, and deliberative system can be seen as debate maps customized so as to be workable in live, problem-solving situations. There are other, more specific differences, but these are largely of implementation detail.¹³

Claims and Argument Classifiers

The scale of a debate map requires that we condense our representation of an argument into a summary title or “claim.” Claims should be brief, informative, and easy to read in context with other claims. In the early rounds of a debate, claims tend to correspond to an argument’s conclusion (“abortion is wrong,” “computers can think”); in later rounds they are more likely to focus on details of reasoning (“computer’s can’t be creative because they follow brittle rules”). In some cases, a claim is simply a well-known title for an argument, as in “The Chinese Room argument.” There is no rule for the formulation of a claim; it is primarily a placeholder which facilitates human interaction with a large debate map. Since claims are placeholders for arguments, I treat “claim” and “argument” synonymously in what follows.

Most claims can be associated with an interrogative—a question which that claim takes a stand on. For example, “Computers can think” corresponds to “Can Computers Think?” and “The Chinese Room Argument” corresponds to “Can Chinese Rooms Think?” Interrogative forms of claims are useful in debate mapping, especially at the level of a “root claim” (the argument to which all subsequent arguments attach), since any region of debate growing from that claim can be thought of as a set of arguments all of which take a position on its corresponding root question.

Debate maps require additional information about arguments and claims. It is useful to know, for a given claim, who made it, in what year, in what journal, from what perspective, etc. Let us call these various items of information “argument classifiers.” Classifiers can be represented by functions from set of arguments in a debate to the relevant set of classifiers (e.g. a map d from arguments to dates, or a map f from arguments to authors). Classifiers are essential to understanding certain debate

structures. For example, by associating arguments with dates, we can note periods in which a debate attracted more or less attention. Or, by associating arguments with “camps” (e.g. Republican vs. Democrat), we can study the unfolding “rounds” in a debate pitting two or more groups against one another. Classifiers are also relevant to the computer implementation of debate-mapping systems, where one might want to highlight all arguments made by a particular author, from the perspective of some theoretical perspective, or in a particular journal or group of journals.

Among argument classifiers, *author* is important but poses special problems. Who to credit as an argument’s author is not always straightforward. In particular, how are we to handle cases where (1) an argument is anticipated, (2) it is articulated by someone besides its original author, or (3) there is no author at all? In the first case, we can designate the argument as being “anticipated by p .” This locates authorship correctly with p but highlights that this is an argument he or she articulates for the express purpose of countering it. In the second case an author p_1 articulates another person p_2 ’s argument (usually with a view to attacking it, and often distorting it), so that it seems unfair to attribute the argument directly to p_2 . In such cases one can designate the argument p_2 as *articulated by p_1* . This correctly locates authorship with p_1 but shows that p_1 intends to capture p_2 ’s reasoning. In the third case, there is *no one* to attribute the argument to. It is just a claim that is (or is assumed to be) actually or potentially held by some group of people. David Lewis (Lewis, 1999), facing this situation, simply calls his interlocutor “Strawman.” For example, p attacks the claim that machines can be creative, or that machines can have free will. But *who* actually says that machines can be creative or have free will? To what author do we attribute this argument which p

disputes? We could try to find someone who says such things, but that is labor intensive and often, the person we find to credit the claim to is *not* someone *p* thought of him or herself as attacking. On the *Can Computers Think?* charts we called such general positions which are widely attacked or supported but rarely explicitly articulated “focus boxes,” insofar as they focus regions of debate by providing attachment points for disputing arguments.

Dispute

I make a claim and you counter. Your argument *disputes* mine; we are now engaged in debate; dialectical structure has begun to emerge.

Dispute and support—our two inter-argument relations—are difficult to precisely define. In the introduction to *Can Computers Think?* we described these as “fuzzy categories” which “cover a wide range of cases” insofar as “claimants respond to one another in a variety of affirmative and negative ways” (p. 8). I propose the following definition of dispute, which covers the cases that emerged in the *Can Computers Think?* project. Barring exceptional circumstances, argument a_2 disputes argument a_1 if the following conditions are met:

- (1) The conclusion of a_2 is the negation of some statement in a_1
- (2) a_2 is relevant to a_1
- (3) The author of a_2 intends for a_2 to dispute a_1

(1)-(3) jointly provide strong inductive evidence for the presence of a dispute relation between two arguments. Consider an example. “Computers can think” is disputed by the argument that “Computers cannot think because they cannot be creative.” The conclusion of the second argument negates the conclusion of the first, the second argument is

relevant to the first insofar as it takes a stand on the question of machine intelligence, and the person who made the second argument intends for it to dispute the claim that computers can think. Thus, the second argument disputes the first. Let us consider each condition on dispute in more detail.

(1) If I develop a counter-argument to your claim, then the conclusion of my argument must negate some statement of yours, at least on suitable paraphrase; otherwise it's hard to see how our two arguments could conflict with each other. Let us call this "d-inconsistency" for dispute inconsistency. That the conclusion of one argument negates some statement in the other is necessary but not sufficient for dispute. A pair of arguments may happen to be d-inconsistent but have nothing to do with one another. We might suppose a simpler condition could be given; namely, that the conclusion of my argument negates the conclusion of yours. But the conclusions of arguments in threads of disputation (discussed below) don't always move from ' p ' to ' $\text{not } p$ ' to ' p ' to ' $\text{not } p$ '... etc. The focus of debate shifts over time. What was an unquestioned premise or enthymeme in an early round of debate becomes an explicit target later on.

(2) Relevance is notoriously difficult to define, even though irrelevance is relatively easy to detect. Fallacies of relevance provide the classic examples. You argue that computers are capable of thinking and I complain about your tone of voice. My complaint is irrelevant to your argument. Or again: you argue that computers can think if they pass the Turing Test. I respond that computers can't think because they cannot be creative. Even though we are both discussing machine intelligence, I have not really responded to your argument.¹⁴

To facilitate relevance detection, one can apply an “interrogative test.” Recall that claims can be converted into questions. To see whether a claim is relevant to the one it disputes, convert the disputed claim into a question and ask whether the disputing claim attempts to answer it. If it does, it is a relevant dispute; otherwise it is irrelevant. In the example just given, the interrogative form of the first claim is the question: “Can computers that pass the Turing Test think?” The counter-claim, “Computer’s can’t think because they cannot be creative” does not take a stand on this issue. Hence the counter is not relevant to its putative target.

(3) The most obvious condition on dispute is intent. One argument disputes another insofar as the author of the first argument says that his or her argument challenges the second one. In creating a debate map one typically begins with intent and then does the work of verifying that the argument fulfills the other conditions on dispute. However, intent alone is not sufficient to establish dispute, as the distinction between surface and deep structure shows. Someone may claim (i.e. have the intent) that his or her argument disputes another one, but the allegedly disputing argument may fail to be either relevant or d-inconsistent with its target, insofar as its author is actually discussing a new point.¹⁵

(1)-(3) are not jointly sufficient for dispute. For example, let the same *reductio ad absurdum* be repeated, and consider whether this counts as a “dispute.” The conclusion of the *reductio* negates one of its own statements, and arguments are trivially relevant to themselves—hence the pair of arguments satisfies (1) and (2). Imagine further that the two arguments are written slightly differently and by different authors, so that the author of the attack does not recognize the target argument as his own repeated.

If the author of the second version of the *reductio* intends for his version of the argument to dispute the first one, (1)-(3) hold even though the argument clearly does not dispute itself.

Even though (1)-(3) are not jointly sufficient for dispute, the satisfaction of all three conditions does provide strong inductive evidence that one argument disputes another. If I am a rational interlocutor in a debate, and I say that my argument disputes yours, then assuming my argument is relevant to and d-inconsistent with yours, it is highly probable that my argument genuinely disputes yours.

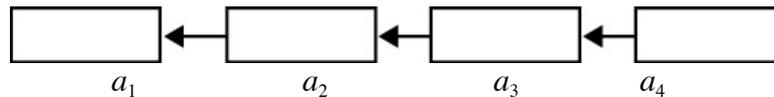
Dispute is irreflexive, asymmetric, and non-transitive; it inherits these properties from condition (3), *intent*.¹⁶ Dispute is irreflexive because rational authors never intend for their arguments to be self-disputing. Dispute is also asymmetric: if I intend for my argument to dispute yours, there is no way you can intend for yours to dispute mine, since yours was framed before mine had even been stated (of course you can refashion your original argument to dispute mine, but that is then a new dispute in the thread). Dispute is non-transitive but not strictly intransitive. Dispute is not transitive insofar as your enemy's enemy is often your friend, or yourself one round earlier. However dispute is not strictly intransitive, since in some cases a person may intend for his or her argument to dispute both a recent argument and an earlier argument in the same thread.¹⁷

Threads

The most basic dispute-based structure is a thread—a sequential chain of arguments which dispute one another. I make a claim. You attack. I counter. And the fight is on. Or, in less combative terms: I make a move. You counter. And the game is

on. Threads are familiar from internet discussion lists, legal proceedings, coffee-shop conversations, congressional activities, and (in impoverished form) day-time talk-shows. We will see in the next section that debates can be viewed as collections of threads which have a shared root argument.

Threads can be represented as directed graphs or “digraphs” whose vertices correspond to arguments and whose edges correspond to dispute relations:



Boxes are used to represent the vertices because, in practice, they contain claims and classifiers. The “argument set” of a thread is the set of arguments within it. Graph—theoretically, threads are finite, acyclic, and connected¹⁸ digraphs with the property that the outdegree and indegree of each vertex is at most one. This means that each argument disputes at most one other argument (outdegree at most one), and is disputed by at most one other argument (indegree at most one). Arguments in the “middle” of a thread dispute and are disputed by exactly one argument. The first argument in a thread (e.g. a_1 above), which disputes no other, is called the “root argument”—it is the basic argument to which all others attach, directly or indirectly. The last, undisputed argument in a thread (e.g. a_4 above) is a “leaf argument;” it represents the final say (at a time) on the subject initiated by the root argument. Arguments in a thread are indexed by natural numbers beginning with the root argument. The “length” of a thread is the cardinality of its argument set, so that the thread above has length 4.¹⁹ A “round” is a pair of arguments a_i, a_{i+1} (where i is typically an odd number), though as we shall see this terminology is more or less apt in different cases. A single argument is a “trivial thread.”

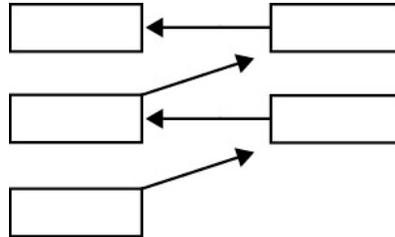
When traversing a thread from root to leaf one proceeds from earlier to later arguments. If the function d maps arguments to dates (and ‘<’ means ‘earlier than’), then in the case above $d(a_1) < d(a_2) < d(a_3) < d(a_4)$. Although dispute relates later arguments to earlier ones (e.g. a_2 disputes a_1), threads are perspicuously displayed as unfolding from left to right, using arrows that go *against* this logical direction (that is, opposite the direction shown in the diagrams above). This leads the eye to track the genesis of a debate, which is useful insofar as early moves in a thread are essential to understanding later, more complex moves. To make clear that such arrows direct the eye but flow opposite the logical direction of dispute, one can label them in the passive voice (e.g., a_1 “is disputed by” a_2).

Threads rarely extend beyond five or six moves. The reason seems to be that as a thread unfolds it becomes increasingly difficult to provide a response which fulfills the relevance condition on dispute. Of course, a conversation will often involve more than five or six exchanges, but by the second or third round of surface debate a discussion will tend to branch off topic. On the *Can Computers Think?* project we called this “thread drift.” The first argument to violate relevance in a thread, the point at which a thread “drifts,” is not for that reason lost. It can be thought of as the root claim of a new thread. In a maximally degenerate conversation, where no party *ever* relevantly addresses the others; we just get a string of unconnected arguments (that is, trivial threads), relevant to themselves only.

To gauge whether an argument is legitimately part of a thread, one needs to ask whether it is relevant to the argument it attacks. If relevance is ensured at each dispositional step, and if relevance is transitive (even though dispute in general is not),

then every argument in the thread should be relevant to its “root claim” a_1 . Threads of arguments which are relevant to one another in this way rarely make it past a fourth round. Likelihood of drift seems to increase the further you are from the root node.²⁰

In the classroom threads can be “crunched together” and depicted as follows:

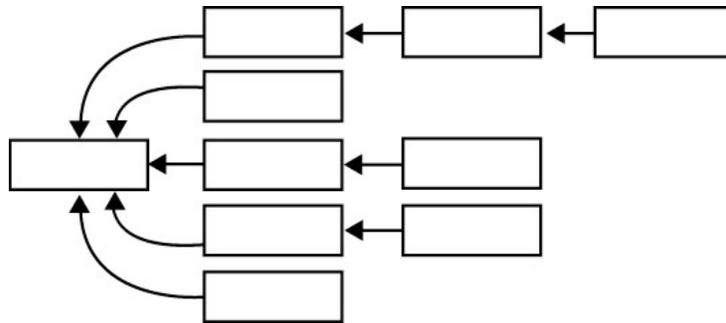


This format makes the “rounds” of a thread obvious, insofar as two columns of arguments are visually “pitted” against one another. Let us call arguments above or below one another in such a column “round separated arguments.” We might expect round-separated arguments to be associated with one or another warring party. Thus, if we associate each argument in a thread with an interlocutor by a function f from arguments to interlocutors, we might expect that $f(a_i) = f(a_{i+2})$ for all arguments a_i, \dots, a_{n-2} in the thread’s argument set. But this is not always the case. Many round-separated arguments are associated with *different* people. One person takes up an argument on someone else’s behalf (say, Brian McLaughlin defending Jerry Fodor against connectionist rebuttals). Moreover, to revisit transitivity, it can happen that one round-separated argument disputes another in the same “column.” This can happen even when round-separated arguments are associated with the same person. It sometimes happens, for example, that in the course of a discussion you modify your view on the basis of previous attacks, even as you continue to resist the attacker.

Debates

I say that Joe should be president. Fifty people respond and I counter them all. Ten die-hard opponents press the fight. Eight of these responses generate counters. I counter two of these but, exhausted, I leave it to my colleagues to carry on the fight. A debate is underway.

Debates can be pictured as follows:



Debates are “rooted digraph trees” where, again, the vertices correspond to arguments and the edges correspond to dispute relations. The “root” argument is the one with outdegree 0, the one that disputes no other. As with threads, the outdegree of each vertex is at most one, but with debates we allow arbitrary indegree: arguments can receive arbitrarily many disputes. The “argument set” of a debate is the set of arguments contained in it. A is a “sub-debate” of B if A and B are debates and the argument set of A is a subset of the argument set of B . A thread is a sub-debate which satisfies the conditions on being a thread. A “maximal thread” is a thread that is not contained in a longer thread. There are as many maximal threads in a debate as there are leaves in its graph; in fact, a debate can be thought of as a collection of maximal threads with a shared root argument. The debate above has 5 maximal threads which overlap at the debate’s

root argument. Maximal threads “overlap” if they intersect at arguments besides their root argument, and a “branch-point” is a non-root argument which is multiply disputed—that is, an argument with outdegree 1 and indegree greater than 1. The argument above contains no branch-points or overlapping maximal threads (though the debate pictured below does).

The assumption that no argument disputes more than one other argument (the outdegree of each vertex is 1 or 0) is *not* plausible, insofar as one frequently directs an argument at multiple targets. On the *Can Computers Think?* project the decision to rule out one-many disputes was made primarily for graphic purposes—it prevents overlapping dispute arrows, cycles, and other graphical features which diminish readability. We addressed the problem in two ways. First, one-many disputes can be visually depicted *without* compromising readability by using hypertext links (which can be printed on paper using small footnote-like pointers). Second, and of more interest, is the fact that multiply-disputing arguments can be separately presented in each of their dialectical contexts and tuned appropriately. This generates, in effect, new arguments, or at least new claims (argument titles written to read well within a thread). In fact, it appears that one generally *does* get a different argument in such cases, and that arguments which are supposed to attack multiple targets in fact differ slightly in each context.

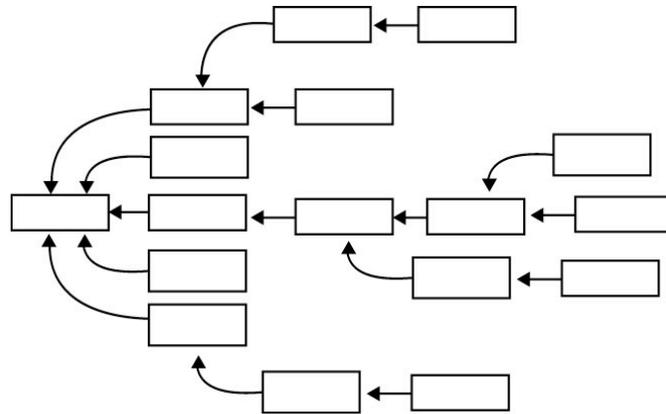
Debates are fruitfully described by a variety of quantities. The “cardinality” of a debate is the cardinality of its argument set. The “length” of a debate is the length of its longest maximal thread. The “spread” of a debate is the indegree of its root argument. The “mean length” of a debate is the mean length of its maximal threads. The “proliferation” of a debate is the number of branchpoints contained in it. For example,

the debate above has cardinality 10, length 4, spread 5, mean length 2.6, and proliferation 0. Again, none of its five maximal threads overlap. None of its threads contains more than two rounds of debate, though whether a “rounds structure” obtains depends on the authors associated with the arguments and the nature of those arguments themselves.

The spread of a debate is a rough measure of the influence of its root argument. Controversial arguments proliferate discussion and attract dispute. The more a debate spreads the more diverse are the arguments that dispute it in a genuinely original way (assuming that redundant arguments are being pruned), and the more difficult it is to spawn new threads.

Whereas thread-length (and hence debate length) rarely extends beyond 5 or 6, debate-spread has no such practical limit (at least in my experience). Debates with spread exceeding 10 or 20 arguments are not unknown (see the articles in *Behavioral and Brain Sciences*, for example), and these tend to generate discussion for years. Perhaps there is some practical upper bound on spread, given that the space of possible answers to a claim will tend to “dry up” over time, causing the focus of the discussion to shift elsewhere. Of course spread is also limited by the fact that consensus can emerge in the course of a debate.

Consider the following, more complex debate:



This debate has cardinality 17, length 5, spread 5, mean thread length 3.3 and proliferation 3. 5 of its 8 maximal threads overlap. No thread contains more than two rounds of debate. I envision a study of debate which uses these and similar quantities to classify debates into structural kinds.

The root claim of a debate is important, because it defines a relevance condition which all subsequent arguments in the debate must fulfill. This implies that in an ideally completed debate-map all cases of thread-drift will have been pruned. Every argument in the debate over John Searle's Chinese Room argument, for example, should be relevant to the Chinese Room argument insofar as it directly or indirectly addresses the question whether or not Chinese Rooms can think.

The "frontiers" of debate are pervasive—any argument in a debate is subject to further disputation. Note the contrast with a literal exploratory map. On a literal map of a partially explored territory, the outermost periphery of darkness constitutes the frontier. But on a debate map the frontiers are everywhere. Every claim is open to further dispute. Of course, the leaves of a debate are, by definition undisputed and open to new argumentation, and the further "inland" you move towards the root claim the more likely

it is that one's objections have already been made by other participants. So debate maps do acquire a kind of explored / unexplored structure corresponding to the distinction between root and leaf arguments and more or less pronounced depending on its length, spread, and proliferation. Nonetheless, in principle the entire argument set of a debate constitutes its frontier.

Structurally, branch-points indicate important sub-debates, since they mark arguments that were controversial enough to provoke a variety of responses. Sometimes sub-debates are referred to by name (e.g. the “imagery debate” relative to the artificial intelligence debate). The interrogative form of the root argument of a sub-debate generally provides a useful name, e.g. “Can Computers Be Conscious?” relative to the sub-debate growing out of the claim that computers can be conscious. Named sub-debates help to organize the many argument on a large debate map, and thereby provide an essential navigation tool.²¹

The techniques described here are applicable to individuals developing their own views, even absent combative partners. Philosophers often develop their views by imagining potential objections and responding to them. Such intra-personal debates can be represented by designating the anticipated arguments “as anticipated by so-and-so” (recall the discussion of authorship above). Whole regions of the *Can Computers Think?* map would sometimes grow out of a single author's article in this way. Debate-structure (spread, width, etc) is in such cases controlled by author.²²

In order to perspicuously grasp the various formal properties of a debate it is useful to embed the debate's graph in a 2-space following the guidelines implicit in the diagrams above: represent threads left to right as growing out of a root claim, and

represent each thread on top of the other. In this way length, spread, and proliferation can be directly grasped. On the (*Can Computers Think?*) project threads were “crunched” together for real-estate purposes. This allowed us to fit the entire debate onto seven charts, but at a cost: the formal structure of the debate (a central gain of the visual format) was largely lost. This can be compensated for by representing debate structures on-line, or by printing only small portions of debates in pages of an “Atlas” (where real-estate is not at such a premium), in such a way as to preserve debate-structure. Historical sidebars, definitions of key terms, and other pertinent information can populate the inevitable white-space.

Support

I make a claim and you support it with a related claim. By marshalling ancillary evidence, citing historical precedent, or introducing a novel argumentation structure, you make an argument distinct from but supportive of mine. Our complimentary arguments are linked by *support*, a second and more complex inter-argument relation. It’s properties are less obvious, its varieties more numerous, its conversational form less perspicuous. My sense is that the concept of support I articulate here will have to be explicitly broken into a family of relations in the course of further investigation. What I offer here are provisional remarks based on my experience with support relations in the *Can Computers Think?* project.

In general, support is used to articulate a position or perspective (e.g. that of the Green party, utilitarianism, creationism, or the gun-control lobby) into a family of mutually complimentary arguments which can serve as separate targets for dispute.

Support as I use it is not (at least currently) envisioned as a mechanism for describing consensus bodies of knowledge.²³ It is envisioned as a means of articulating a body of knowledge *relative to a debate*. Without support relations, it is very difficult to contextualize a set of arguments; one ends up with free-floating threads and debates which cannot be organized into a larger structure. By providing a connected set of attachment points for disputing arguments, supports help structure a large-scale debate in a coherent way.

On the *Can Computers Think?* project we made use of at least three separate concepts of support: (1) logical, (2) historical, and (3) specialization. What unifies these three varieties of support is, again, that they can be used to articulate a body of knowledge into distinct targets for attack. In each case following I describe a version of support and a necessary (but not sufficient) condition on it. Since relevance and intent function similarly in this context to the way they do in the context of dispute (with the exception of intent in the context of historical support, as we shall see), they are not discussed further below. For any pair of arguments if relevance, intent, and the listed necessary condition obtain, we have strong inductive evidence that that type of support obtains.

Logical support. Argument a_2 logically supports argument a_1 if the conclusion of a_2 strengthens the conclusion of a_1 . Arguments which logically support one another can be thought of as parts of a complex argument. Suppose a_2 supports a_1 insofar as the conclusion of a_2 adds to the inductive strength of or figures in a deduction to the conclusion of a_1 . Then argument a_2 can be thought of as a sub-argument of the argument which results if we join a_2 and a_1 . Logical support is often provided by specific

experimental or empirical results. For example, the claim that “Computers can understand analogies” is supported by various computer programs which demonstrate analogical reasoning skills. Each of these analogical reasoning programs is subject to individual critique, as is the general claim that computers understand analogies.

Historical or context support. Argument a_2 contextually supports argument a_1 if it is an earlier argument or contextually relevant argument which the person who makes a_1 draws on. Context supports are important because they situate an argument (and thus a debate or position) in its wider disciplinary and historical context. For example, many of Hubert Dreyfus’ arguments in the artificial intelligence debate are supported by the earlier phenomenological claims of Martin Heidegger. Note that with contextual support intention works in the opposite direction as with dispute and other forms of support: one author generally intends to support him or herself with the historical or contextual arguments of another, rather than intending that his argument historically or contextually support another.

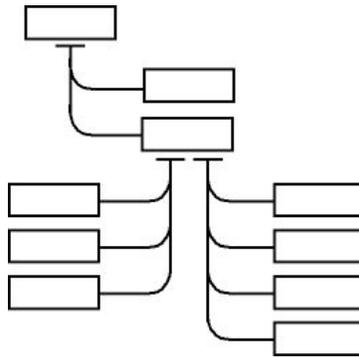
Specialization. Argument a_2 supports argument a_1 through specialization if it is a more specific version of a_1 . Consider the case where I argue in favor of machine thinking with reference to features of connectionist networks (a specific kind of computational architecture in artificial intelligence). Then the claim that connectionist networks can think is a special case of the claim that machines think, and is supported by it in that sense. Specialization supports also correspond to specific versions of a general claim (typically in a focus box) made by individual authors. For example, the general claim that “Consciousness arises from higher-order representational structures” is

supported by specific versions of that claim in Johnson-Laird, Dennett, Rosenthal, and Pollack.

As with dispute, support is irreflexive, asymmetric, and non-transitive (though it is represented as strictly intransitive in its visual representation), and again it inherits these properties from intent.

Positions

We have seen how arguments related by dispute form into threads and debates. Similarly with support. Arguments related by support form into thread-like sequences which I call “support paths.” These in turn aggregate into “positions.” Support paths and positions have the same graph-theoretical representation as threads and debates, respectively; they only differ in the interpretation given to the edges. Following is a position with 10 arguments and 8 support paths:



Examples of positions include pro-Choice, pro-Life, “materialism” in the philosophy of mind, the canons of the “republican” or “democratic” parties, and “utilitarianism” in ethics. An argument by itself is not only a trivial debate, it is also a trivial position.

We might imagine that the diagram above is part of an ethics debate map, and that the highest level argument is a utilitarian focus box, supported by act and rule-

utilitarianism, the latter supported by eight specific versions of rule-utilitarianism associated with specific authors. Disputers can then direct their energies either at utilitarianism wholesale or at specific versions of rule-utilitarianism.

Graph theoretically, support paths are identical to threads and debates are identical to positions—it is simply the interpretation of the edges which differs. Of course, the fact that support may proliferate into three varieties complicates matters. A path of logical supports is a kind of complex argument. Paths of historical supports correspond to historical threads which follow from one claim to another backwards through time. Paths of specialization correspond to sequences of increasingly specific versions of an argument. Collections of such paths correspond to hierarchically organized taxonomies which describe the various ways a particular argument ramifies.

In each case we could define length, spread, mean path length, proliferation, etc., in the same way, though perhaps other terms would be better suited. Further study is called for, and for now I leave it open which measures would be most useful in a logic of support.

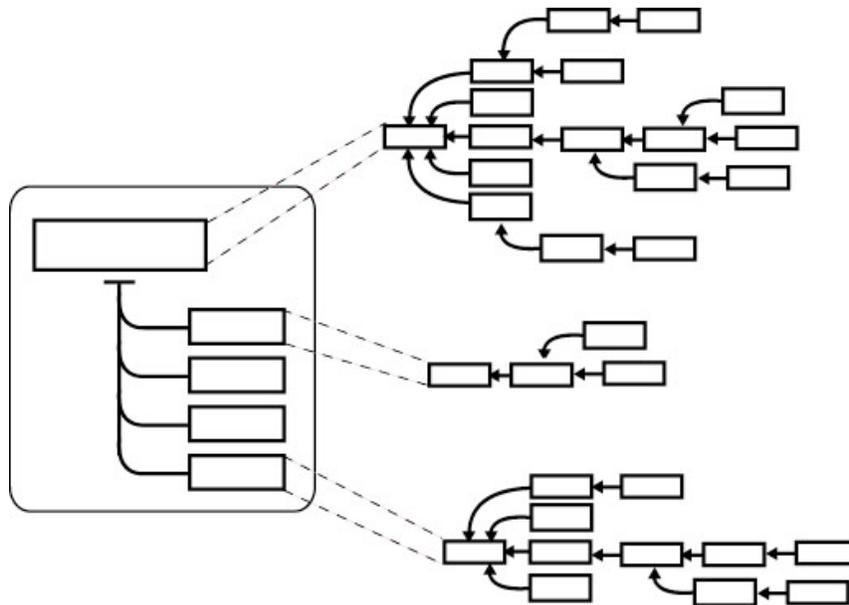
As with dispute, support is perspicuously depicted as flowing opposite its logical direction, using the passive label “is supported by.” This is not to force us to move diachronically from earlier to later arguments, but so we stay focused on the primary conceptual structure of an argument and only move to historical context, ancillary evidence, or a more specific version if we are so inclined.

Supporting links are, for obvious reasons, best positioned *below* the arguments they support. A “t-bar” symbol is adopted for contrast with the dispute arrow and to further capture the “support” metaphor. It is left to future versions of debate maps to

introduce graphics conventions for distinguishing different forms of support. Also note that, as with dispute, one-many support relations were not used on the charts, in order to prevent visual spaghetti. Again, this can be compensated for by hypertext links. And again, the process of rewriting a supporting argument for some particular context typically involves changing its content enough that it becomes, in effect, a new argument.

Combined Structures

The picture that emerges is one where large sets of arguments aggregate into positions and debates in virtue of support and dispute relations. There are a number of ways debates and positions can interact. One, large scale form of interaction is best expressed graphically, by making positions the focus and representing “clouds of debate” attaching to these:



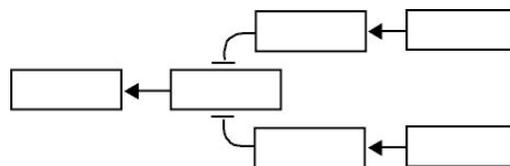
The zoom-out lines indicate that the same argument has been repeated so that it can be seen both in the context of a position and as the root claim in a debate. Above we have a

position comprised by five arguments and four supports, to which three non-trivial debates attach. If the supports above are specialization supports, then two debates attach to specific versions of the general claim at the top of the position, and the largest debate targets the general claim itself.

The intersection of the argument set of a position and the union of the argument sets of a collection of debates is a “front”—the position’s interface with its dialectical context, that subset of arguments within it which have been critically engaged. Suppose we posit a set A of published arguments, a set $D \subset A$ of published arguments which are situated in non-trivial debates, and a set $P \subset A$ of arguments situated in non-trivial positions. Then, the intersection $D \cap P$ constitutes a massive battle front between articulated belief systems (codified in published arguments) and critical engagements with the same.

This is just one of several ways positions and debates can interact. It is also possible that positions could spawn debates which could spawn positions in a kind of chain. Or it could be that positions are what are basic and that when they “rub up” against one another they spark debates, and that the insights generated by those debates are incorporated over time into the two positions, causing them to become more subtle and complex via a kind of large-scale dialectical process.

Another use of positions is as a supplement to the representation of debates. Sometimes an argument can be represented as a small position, each of whose arguments is separately disputed:



Note that the supporting arguments are represented both above and below the argument they support, in order to preserve the graphical “flow” of the thread structure, I leave formal analyses of this kind of “complex thread” for separate study, though one likely approach would be to treat mutually supporting arguments (positions embedded in threads) as single arguments relative to the debate, so that the same quantities as are used in analyzing a simple thread can be used to analyze the complex thread. In fact, if the support in question is a logical support, then one can either treat the position as a set of arguments linked by supports or as a single complex argument which is disputed by arguments targeting different statements within it (whether to expand a complex argument into a small position or leave it as a single argument is another area where the mapmaker’s judgment is required).

Conclusion

If there is no generally accepted answer to the question posed by a debate map (the interrogative form of its root argument), then it charts a live debate, and serves as an aid to reasoning. Newcomers to the debate are prevented from repeating old moves, and are directed to promising regions of thought. On the other hand, if consensus has developed on the map’s root question, then it is a historical document which charts the critical history of a piece of knowledge, the sequence of moves which had to unfold before the relevant claims were generally accepted. If consensus develops on the question concerning machine intelligence, then the “Can Computers Think?” maps will become historical in this sense. But since the question remains open, these maps

currently serve as a deliberative aid to philosophers, cognitive scientists, and other participants in the debate.

In addition to serving heuristic and documentary needs, debate maps also shed light on the properties of large-scale debates. Using the tools introduced above we can ask, for example: What is the average spread of a policy debate? How often is relevancy violated in a presidential debate? What percentage of arguments in the abortion debate are named, and what percentage are “focus boxes”? As more data is gathered from more debates, general properties of debate—perhaps even laws—may emerge.

An upshot of such studies could be more fruitful public and private debate, insofar as awareness of the deep structure of debate could promote better practice in producing its surface structure.

Author's Note

I am grateful to Mark Deering, Jean Goodwin, Robert Horn, Scott Hotton, Rolf Johansson, Paul Livingston, P.D. Magnus, Wayne Martin and an anonymous reviewer for helpful comments and critique.

References

- Antonio, A. L., Chang, M. J., Hakuta, K., Kenny, D. A., Levin, S., Milem, J. F., et al. (2004). Effects of Racial Diversity on Complex Thinking in College Students. *Psychological Science, 15*(8), 507-510.
- Beardsley, M. C. (1950). *Practical logic*. Englewood Cliffs, N.J.: Prentice-Hall.
- Conklin, J. (2003). Dialog Mapping: Reflections on an Industrial Strength Case Study. In *Visualizing argumentation : software tools for collaborative and educational sense-making*. London ; New York: Springer.
- Eagleman, D. M., & Holcombe, A. O. (2003). Improving science through online commentary. *Nature, 423*(15).
- Eemeren, F. H. v., & Grootendorst, R. (1992). *Argumentation, communication, and fallacies : a pragma-dialectical perspective*. Hillsdale, N.J.: L. Erlbaum.

- Freeley, A. J., & Steinberg, D. L. (2000). *Argumentation and debate : critical thinking for reasoned decision making* (10th ed. ed.). Belmont, CA: Wadsworth/Thomson Learning.
- Freeman, J. B. (1991). *Dialectics and the macrostructure of arguments : a theory of argument structure*. Berlin ; New York: Foris Publications.
- Goodwin, J. (2001). Wigmore's Chart Method. *Informal Logic*, 20, 223-243.
- Horn, R.E., Yoshimi, J, Deering, M, McBride, R. (1998). *Can Computers Think*. MacroVU, Inc.
- Horn, R. E. (2003). Infrastructure for navigating interdisciplinary debates: critical decisions for representing argumentation. In *Visualizing argumentation : software tools for collaborative and educational sense-making*. London ; New York: Springer.
- Kirschner, P. A., Buckingham Shum, S. J., & Carr, C. S. (2003). *Visualizing argumentation : software tools for collaborative and educational sense-making*. London ; New York: Springer.
- Lewis, D. K. (1999). Reduction in Mind. In *Papers in metaphysics and epistemology* (pp. ix, 453 p.). Cambridge, UK ; New York : Cambridge University Press.
- Rescher, N. (1977). *Dialectics : a controversy-oriented approach to the theory of knowledge*. Albany: State University of New York Press.
- Toulmin, S. E. (2003). *The uses of argument* (Updated ed.). Cambridge, U.K. ; New York: Cambridge University Press.
- Van Gelder, T. (2003). Enhancing Deliberation Through Computer-Supported Argument Visualization. In *Visualizing argumentation : software tools for collaborative and educational sense-making*. London ; New York: Springer.

¹ Multiple parties are not required. Internal dialogues can have the structure of a debate. Rescher (1977, ch. 3) calls this “unilateral dialectics.”

² See (Rescher, 1977), pp 1-3 for historical background.

³ See (Eemeren & Grootendorst, 1992), p. 3 for background.

⁴ See (Kirschner, Buckingham Shum, & Carr, 2003) for historical and contemporary references.

⁵ For Freeman, "microstructure" involves the "internal structure of statements" (p. xi). Relative to this, his study of relations between statements in arguments is indeed "macrostructure." However, I focus on an even higher level macro-structure, which treats whole arguments as basic elements.

⁶ See (Freeman, 1991), chapter 1, for historical background on argument diagrams.

⁷ Freeman traces the “standard approach” to argument diagramming to (Beardsley, 1950).

⁸ Versions of a debate map in this sense exist, most prominently in the “flow sheets” and other diagrams used in judging competitive debates (See Freeley & Steinberg, ch. 14), but also in theoretical contexts (see Rescher, p. 5, p. 17).

⁹ The notion of a transform from surface to deep structure can be made especially perspicuous by analysis of a structured debate. In a structured debate, a predetermined group of proponents, opponents, and judges engage in debate according to a set of rules. All structured debates of a given kind have, at least in principle, the same surface structure, and will thus be represented by isomorphic debate maps. Nonetheless, the deep structure is likely to vary from case to case. Thus 10 different forensic debates on the same topic, carried out according to the “cross-examination format” (Freeley & Steinberg, ch. 18) will each manifest a surface structure consisting of twelve moves arranged on a single thread, but are likely to vary considerably in their deep structure.

¹⁰ This is clearly just an ideal. As noted below, the construction of a debate map involves interpretation on the map-maker's part, and this interpretive dimension is more pronounced in certain contexts (e.g. determination of "relevance").

¹¹ If a participant in a debate retracts a claim which is held by others, it can remain on the map, especially if others have responded to it explicitly or used it to support their own arguments.

¹² Insofar as the goal on such maps is not to "win an debate," but rather to genuinely understand an issue and come to some resolution on it. Hence the live deliberative process mapped in such cases attempt to proceed along disputational lines which involve links to relevant and non-redundant arguments.

¹³ (1) In deliberative systems what are related are single statements, not whole arguments. However, what are essentially being used are what I call "claims" (argument titles), and in general any node on a deliberation map *can* be associated with a full argument to that conclusion. (2) The graphic conventions followed by Conklin and Van Gelder are somewhat different than those I endorse, though since the underlying graph theory is identical this is obviously not a fundamental difference. Van Gelder introduces some particularly effective visual conventions. (3) In the IBIS system (Conklin, 2003) questions are treated as basic elements. Thus one moves from questions to proposed solutions to arguments, etc. In what follows I show that any claim (argument title) can be associated with an interrogative form (and conversely, any interrogative can be associated with a particular answer to it), so that the two systems are interchangeable in principle.

¹⁴ Of course, relevance is still relatively ill-defined. It may be that inter-argument relations are best characterized by *degrees* of relevance. Note moreover that relevance is an area where mapmaker interpretation is prominent.

¹⁵ In some cases argument a_2 clearly disputes a_1 though the author of a_2 does not explicitly mention the author of a_1 . Such disputes are natural to include in a map, but require more interpretation on the mapmaker's part. On the Can Computers Think? project we distinguished "stated" from "implied" dispute in exactly this way. Implied dispute will not be considered in this paper.

¹⁶ Note that relevance is reflexive and may be transitive as well. This has the result that later arguments should be relevant to all previous arguments in a "thread," as we see in the next section.

¹⁷ Disputes are visually represented as strictly intransitive (no dispute links "leap frogging" one another). As we shall see, this is done both for visual clarity and to enable us to distinguish the same argument as it works against multiple targets.

¹⁸ All are *prima facie* plausible. Threads cannot go on forever, at least in practice, hence finitude. Acyclicity is required insofar as cyclicity would require that a later argument could dispute an earlier argument. A collection of arguments which is not connected is best thought of as a set of threads, rather than a single thread, hence connectedness.

¹⁹ I here overload the term "length," insofar as the length of a path in graph theory corresponds to the number of edges in that path. I use it to refer to the number of vertices (arguments) in a path.

²⁰ It is interesting to note that forensic debate formats typically mandate 4-7 speaking turns (see Freeley & Steinberg, ch. 18), which result in 8-14 node threads of surface structure which are likely, based on this analysis, *not* to correspond to single threads of deep structure (multiple threads may emerge, arguments may be repeated, etc.).

²¹ We imposed a three-tiered regional hierarchy on the *Can Computers Think?* project. The highest tier was the debate as a whole, the entire set of 7 charts. The root argument for this debate was Turing's (1950) claim, that machines can or will be able to think. Every argument on the entire set of charts, insofar as thread-drift was enforced, directly or indirectly weighs in on the issue of machine-thinking. Second were the physical charts themselves, each with its own root claim and interrogative label (e.g. Can Chinese Rooms Think? Can Connectionist Networks Think? Are Thinking Machines Mathematically Possible?). The lowest tier corresponded to what we called "regions," sub-debates with 2-30 arguments.

²² Rescher develops a "disputational model of inquiry" on this basis. See (Rescher, 1977), ch. 3.

²³ Sets of statements which are mutually supportive are of course well studied. Theories, for example, are sometimes understood as sets of statements closed under entailment.