

Scientific argumentation detection as limited-domain intention recognition

Simone Teufel

Computer Laboratory
University of Cambridge
15 JJ Thomson Avenue, Cambridge, UK
Simone.Teufel@cl.cam.ac.uk

Abstract

We describe the task of intention-based text understanding for scientific argumentation. The model of scientific argumentation presented here is based on the recognition of 28 concrete rhetorical moves in text. These moves can in turn be associated with higher-level intentions. The intentions we aim to model operate in the limited domain of scientific argumentation and justification; it is the limitation of the domain which makes our intentions predictable and enumerable, unlike general intentions.

We explain how rhetorical moves relate to higher-level intentions. We also discuss work in progress towards a corpus annotated with limited-domain intentions, and speculate about the design of an automatic recognition system, for which many components already exist today.

1 Introduction

Automatically recognising the structure of an argument is an attractive and challenging task, which has received interest for a long time from the AI as well as the natural language processing community, and recently from both communities together in a joint effort. Because arguments are global text structuring devices, argument recognition has the potential to advance text understanding and the many real-life tasks that could profit from it.

There are various definitions of what an argument is (Toulmin, 1958; Cohen, 1984; Dung, 1995; Brüninghaus and Ashley, 2005; Besnard and Hunter,

2008; Walton et al., 2008; Green, 2014). We are here interested in a definition close to discourse structure, and concentrate in particular on the recognition of prototypical argumentation steps in scientific exposition. We posit that these argumentation steps can be defined at an abstract level so that world knowledge is not required for their recognition.

There is a clear connection between our goal and intention recognition. Fully understanding every aspect of an author’s argumentation requires the recognition of all of their intentions, which in turn means that we would have to model, generalise over, and do inference with general world knowledge. This is of course an AI-hard task fraught with many theoretical and practical problems; consider the symbolic AI work on this and closely related problems (e.g., Schank and Abelson, 1977; Pollack, 1986, 1990; Norvig, 1989; Cohen et al., 1990 and Carberry, 1990).

We will propose instead to reframe argumentation detection as a *limited-domain* intention recognition task. The basic building blocks of our model of an argument are instances of higher-level intentions which the authors are likely to have had when they were writing their paper. The representation we suggest for intentions does not contain any propositional content based on arbitrary world knowledge. Instead, our intentions are represented as generalised propositions such as “Our solution is better than the competition’s”. Such speech acts realise parts of the author’s intention of persuading the reader that the work described in the paper is novel and significant. When during processing we encounter the sentence

To our knowledge, our system is the first one

aimed at building semantic lexicons from raw text without using any additional semantic knowledge.

(9706013, S-171)

our representation only registers the author's intention of staking a novelty claim for their new work. The proposition is generalised in that the propositional content of the novelty, i.e., the fact that the authors built the first lexicon from raw text without any additional semantic knowledge, is not encoded. This detail is not important at the level of abstraction we have in mind.

The simplification of argument recognition into a limited-domain intention recognition problem is possible because of the high degree of conventionalisation of scientific argumentation. Following Swales (1990), we call explicit statements such as the above novelty claim "rhetorical moves". Rhetorical moves are well-documented in various disciplines: they occur frequently, and they can be enumerated and classified, as applied linguists have done in some detail for several disciplines (e.g., Myers, 1992; Hyland, 1998; Salager-Meyer, 1992).

Swales also coined the expression "research space" – a cognitive construct consisting of scientific problems, methods and research acts that authors use when they locate their research with respect to historical approaches and current trends.

When we faced the decision of which types of semantic participants to encode in our representation of rhetorical moves, we tried to achieve as much generalisation as possible, in line with the Knowledge Claim Discourse Model (KCDM, Teufel, 2010). In fact, the core semantic participants in rhetorical moves can be reduced to just two sets – US (the paper's authors) and THEM (everybody else who has ever published).

When it comes to the states and events expressed in rhetorical moves, we maximally generalise again and end up with four classes of predicates, where the classes are defined based on the number of participants in the logical act expressed in the move. We differentiate statements about the authors' own work (US); statements about others' previous work (THEM); statements about the connection between the authors' work with previous work (US and THEM); and finally statements about the research space and the authors' position in it. Another relevant observation is that rhetorical moves often con-

tain sentiment, in the form of "good" vs. "bad" situations, as well as successful vs. failed problem solving acts.

As far as the representation of time in the events and states described in rhetorical moves is concerned, another simplification is possible: it suffices to model three points in time, the time before the authors' research activity begins (t_0), and the times during (t_1) and after (t_2) their research activity. Of course, the real actions by the authors that gave rise to the research in the paper are spread in time in far more complex ways, but a scientific paper is a social construct (Bazerman, 1985). The telling of "the story" follows the convention that all research acts associated with the paper happen simultaneously, and that they transform an earlier state of the world into a new (better) one.

These simplifications allow us to define the 28 rhetorical moves in Figure 1¹. We also give some examples of rhetorical moves from the chemistry, computational linguistics and agriculture literature, which were sourced from our annotated corpora.

The overall argumentation structure we propose concerns the author's argument that their research was worthy of publication, and all of its subarguments – which, at its heart, is always the same argument. Argument recognition then corresponds to a guess as to which strategy the author pursued in making this argument. This process will have to be driven by a bottom-up recognition of rhetorical moves, as these are the only explicitly expressed parts of the argument. This will trigger a simple form of inference as to which higher-level intention might have been present during the writing of the paper.

In previous work, we have used a robust classification model called Argumentative Zoning (AZ; Teufel, 2000, 2010; Teufel et al. 2009, O'Seaghdha and Teufel 2014), that turns some aspects of the more general argumentation recognition model of the KCDM into a simple sentence classification task. In AZ, rhetorical moves with a similar function were bundled together into 7 (in later versions 15 or 6) flat classes or zones, and each sentence was classified into one of these on the basis of surface features,

¹An earlier version of the list of moves appears in Teufel (1998).

I. Properties of research space	
R-1	Problem addressed is a problem
R-2	New goal/problem is new
R-3	New goal/problem is hard
R-4	New goal/problem is important/interesting
R-5	Solution to new problem is desirable
R-6	No solution to new problem exists
II. Properties of new solution (US)	
R-7	New solution solves problem
R-8	New solution avoids problems
R-9	New solution necessary to achieve goal
R-10	New solution is advantageous
R-11	New solution has limitations
R-12	Future work follows from new solution
III. Properties of existing solution (THEM)	
H-1	Existing solution is flawed
H-2	Existing solution does not solve problem
H-3	Existing solution introduces new problem
H-4	Existing solution solves problem
H-5	Existing solution is advantageous
IV. Relationships between existing and new solutions (US and THEM)	
H-6	New solution is better than existing solution
H-7	New solution avoids problems (when existing does not)
H-8	New goal/problem/solution is different from existing
H-9	New goal/problem is harder than existing goal/problem
H-10	New result is different from existing result
H-11	New claim is different from/clashes with existing claim
H-12	Agreement/support between existing and new claim
H-13	Existing solution provides basis for new solution
H-14	Existing solution provides part of new solution
H-15	Existing solution (adapted) provides part of new solution
H-16	Existing solution is similar to new solution
Recently, R-4 <i>the use of imines as starting materials in the synthesis of nitrogen-containing compounds has attracted a lot of interest from synthetic chemists.</i> ⁽¹⁾ (b200198e)	
H-4 <i>This account makes reasonably good empirical predictions, though H-2 it does fail for the following examples: ...</i> (9503014, S-75)	
H-12 <i>Greater survival of tillers under irrigated conditions agrees with other reports in barley [4,28] and wheat [10,13,26].</i> (A027)	

Figure 1: Rhetorical moves; some examples

including sequence information. This way of phrasing the problem allows for tractable recognition and evaluation. AZ classification has been shown to lead to stable and reliable annotation on several scientific disciplines, and it is also demonstrably useful for a set of applications such as the detection of new ideas in a large scientific area, summarisation, search, and writing assistance.

Nevertheless, AZ is only a flat approximation of a larger argumentation model of scientific justification. The work presented here is a departure from AZ in that it aims to model the stages of scientific argumentation in a more informative, finer-grained way.

2 The role of citations in the argument

The reader may have noticed that the rhetorical moves in parts III and IV of Fig. 1, which are concerned with statements about THEM (i.e., other published authors), are closely connected to citation function². In fact, we have in the past attempted the recognition of some of the H-moves as an isolated task, in the form of citation function classification (CFC; Teufel et al., 2006); others (Garzone and Mercer, 2000; Cohen et al., 2006) have used other schemes for similar citation classification tasks.

Where, how often, and how authors cite previous work is an important aspect of their overall scientific argument. For instance, the authors might choose one of the possible articles types (review, research paper, pioneer work etc) to support a particular point in their overall argument. The choice of a particular pioneer paper might signal their intellectual heritage. They might tell us who their rivals are, and who uses similar methods for a different goal (i.e., not rivals), whose infrastructure they borrow, and whose work supports theirs and vice versa. These questions will crucially influence where in the text (physically and logically in terms of the argumentation) a given citation will occur.

As a result of all this, it is often possible to determine some citations as being particularly central to the authors' paper. This information, if it could be automatically determined from text in a reliable

²These 16 moves also follow a different naming scheme, where the move name starts with the letter "H" – historically, such moves were called "hinge" moves, as opposed to the "R" ("rhetorical") moves in parts I and II of Fig 1.

way, would vastly improve bibliographic search. It also has the potential to improve bibliometric assessments of a piece of work's impact, e.g. in the sense of Borgman and Furner (2002), White (2004), and Boyak and Klavans (2010).

3 Higher-level intentions

There are some rhetorical moves that at first glance seem to make little sense. Stating H-5, praise of other people's work, might comparatively weaken the author's own knowledge claim. Similarly, stating H-9, the fact that the author's research goal is harder than other people's goal, might prompt the criticism that the authors have simply chosen their goal badly – had they chosen an easier goal, the solution might have been easier, or achieved better results.

However, rhetorical moves must be interpreted as part of the larger picture of the overall scientific argument. Scientific writing can be seen as one big game where an author's overall goal is to successfully manoeuvre their paper past the peer review, so that it can be published.

According to the conventions of peer review, there is a small set of criteria for acceptance – the authors need to show that the problem they address is justified (High-Level-Goal 1 or HLG-1 for short), that their knowledge claim is significant (HLG-2) and novel (HLG-3), and that the research methodology they use is sound (HLG-4). If valid evidence for the fulfilment of these criteria is presented, the peer review cannot justifiably reject the paper.

Fig. 2 spells out how the overall argument for validity is put together from high- and medium-level intentions and rhetorical moves³. Rhetorical moves in Fig. 2 appear in shaded boxes (H- and R-type moves in different shades of grey). Above the rhetorical moves, we see a simple representation of the intentions posited in the model. For simplicity and readability, Fig. 3 repeats the same network without rhetorical moves. The arrows in both figures express the “supports” relationship in argumentation theory. For instance, in order to argue for the novelty of one's work, a state-of-the-art comparison may or may not be necessary – this depends on whether one describes the research goal as new or not. For new

³An earlier version of this diagram appears as Fig.3.1.7 in Teufel (2000, p.105).

research goals, one may simply show that no other work is similar enough to one's goal: new goals (created at t_1) cannot be compared to existing state-of-the-art, which is frozen in time at t_0 . (Novelty is a rare example of a high-level intention which can be left to the reader to infer, or alternatively stated explicitly as move R-2 or R-6.)

Note that each citation that has an H-type rhetorical move associated with it automatically strengthens the claim that the authors are knowledgeable in the field (one of the important subgoals of HLG-4, soundness). Under our model, citations without any associated H-move are not contributing to this goal, as a knowledgeable author must be able to state the relationship of the current work to earlier work. (A simple statement of similarity with somebody else's work should barely count, but has been given a “weak” move, H-16, because we encountered it so frequently in our corpus studies.)

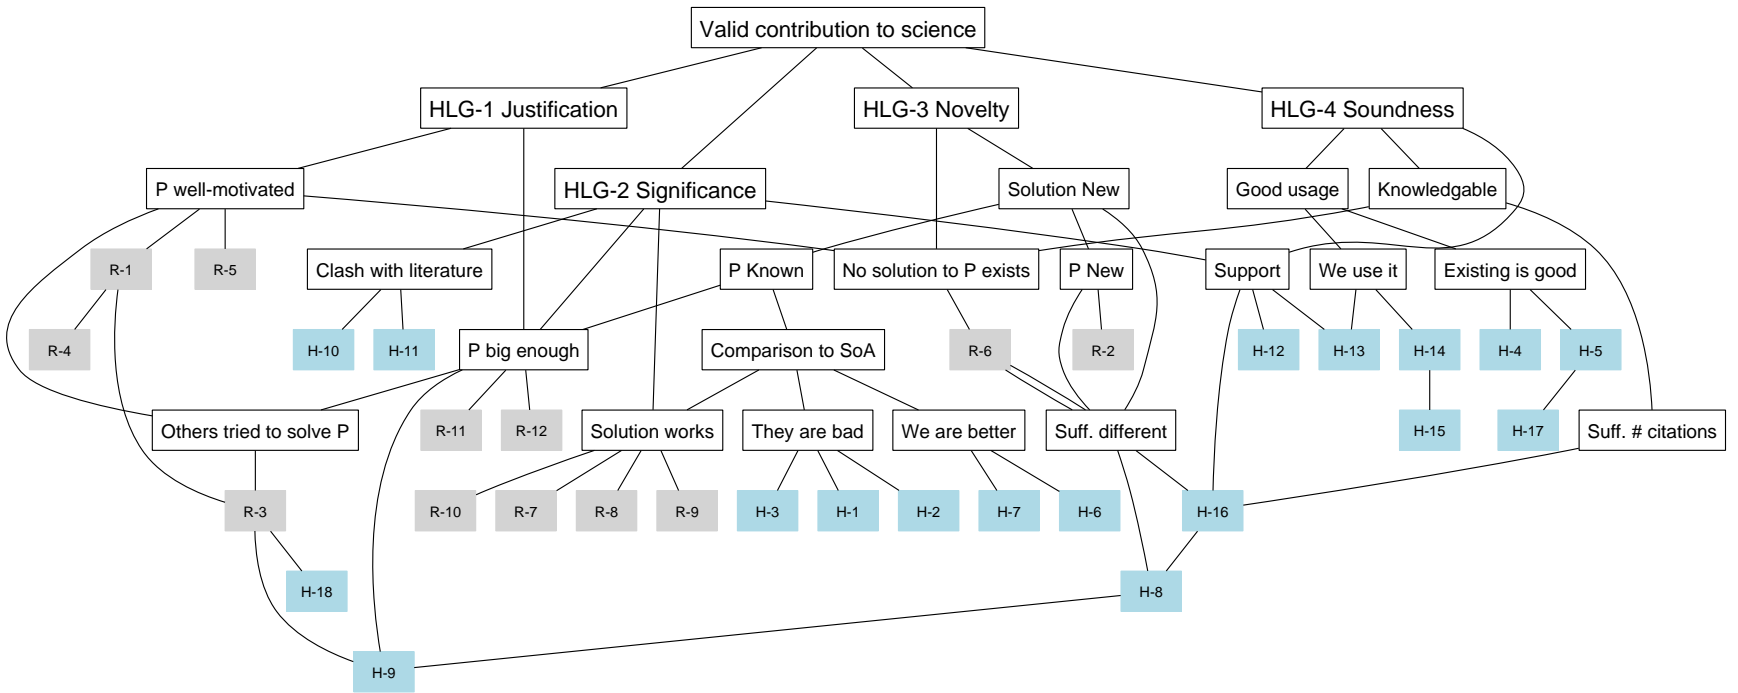
From Fig. 2 we can now see why stating H-5 can be a good strategic move even though it praises other people's work – it supports HLG-4 (soundness of methodology) via the sub-argument that by including praise-worthy existing work, the authors make sure they use the best methods currently available. Similarly, the statement that one's goal is harder than somebody else's motivates that the authors' chosen problem is justified (HLG-1) and significant (HLG-2), and additionally strengthens HLG-4 (via the claim that the authors know their field well). This illustrates that a rhetorical move can support more than one high-level intention.

4 Knowledge representation of moves and intentions

What has been said so far raises the question of which knowledge representation is most suited for modelling intentions and rhetorical moves. Designing a propositional logic that expresses the full semantics of rhetorical moves and of higher-level intentions is a task that goes far beyond the current paper; it requires a thorough design of the semantics of objects and events/states in this limited domain, as well as an appropriate type of inference. Nevertheless, we will sketch some of the principles of what might be usefully encoded.

The THEM entities would need to be grounded to

Figure 2: Argumentation network (including rhetorical moves).



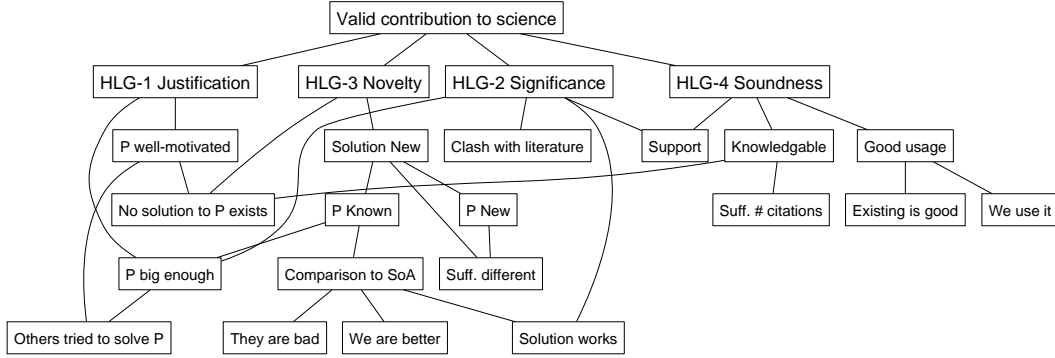


Figure 3: Argumentation network (excluding rhetorical moves).

citations, possibly also to more general entities such as “many linguists in the 1970s”. Entities would need to be tracked throughout the paper, for instance by performing co-reference. We would also need to represent problems, solutions and goals as atomic types, i.e., the fact that they are considered problems, solutions and goals, rather than their content. (The system should keep pointers to the textual strings that express this content, so that down-stream processing or human users can gain access to this information.)

The exact representation of a proposition is open to speculation at this point, but moves would likely be decomposed into atomic clauses. Events and properties in the limited domain (such as changing a solution into another one, or the fact that one solution is better than another) would be associated with a time; for instance all actions that logically happen during the research act presented in the paper would be associated with t_1 .

Inference could be performed by a theorem prover, which could inhibit or further activate the potentially possible “supports” relationships given in Fig. 1, by taking the plausibility of a particular inference into account, in the light of the textual evidence encountered.

Axioms could directly encode some of the rules of the scientific publication game, such that the existence of a problem is a bad state, that of a solution is a good state, but that a solution *needing* something else is a bad state again. Temporal inference could require axioms such as things that persist at a certain time also persist in later times, unless they are changed.

R-5	$\text{solution}(s) \wedge \text{solve}(s, p, t_1) \wedge \text{good}(a, t_2) \wedge \text{aspect}(a, s) \wedge \text{problem}(p) \wedge \text{address}(\text{US}, p)$
R-12	$\text{problem}(p_1) \wedge \text{cause}(s, p_1, t_1) \wedge \text{solution}(s) \wedge \text{solve}(s, p) \wedge \text{problem}(p) \wedge \text{address}(\text{US}, p)$
H-1	$\text{solution}(s_1) \wedge \text{own}(\text{THEM}, s_1) \wedge \text{bad}(a, t_0) \wedge \text{aspect}(a, s) \wedge \text{solve}(s_1, p) \wedge \text{problem}(p) \wedge \text{address}(\text{US}, p)$
H-7	$\text{solution}(s_1) \wedge \text{own}(\text{THEM}, s_1) \wedge \text{solution}(s) \wedge \text{own}(\text{US}, s) \wedge \neg \text{solve}(s_1, p, t_0) \wedge \text{solves}(s, p, t_1)$
H-15	$\text{own}(\text{THEM}, s_1) \wedge \text{solution}(s_1) \wedge \text{solution}(s_2) \wedge \text{change}(\text{US}, s_1, s_2, t_1) \wedge \text{use}(\text{US}, s_2, t_1)$

Figure 4: Sketch of knowledge representation for selected rhetorical moves

As an example of what the representation might look like, Fig. 4 expresses five moves in a simple propositional logic. Here, ownership of solutions (by US or THEM) is expressed directly, as are simple relationships between solutions, problems, results and claims. Consider move H-15, for instance – adapting somebody else’s solution means taking it, changing it into something else, and then using the changed solution. Some moves, such as R-6 and R-9, look like they might require quantification, which exceeds the expressivity of simple predicate logic.

Several aspects of the moves’ semantics are not explicitly expressed in text; they could even be modelled as presuppositions. For instance, R-7 states that a rival’s solution does not solve one’s problem, which presupposes that the author’s solution does, otherwise it would not be a relevant statement. R-7 thereby implicitly invokes a comparison between the

author's approach and the rivals', which is won by the authors. Crucially, whether or not the authors' successful problem-solving is explicitly mentioned in the text or not is optional. Another example is the need to know whether a problem mentioned in a certain rhetorical move is actually the problem that the authors address in the current paper. This is often decisive, because the knowledge claim of the paper is connected exclusively to this particular problem. In some part of the paper, the authors give us the information which problem it is that they address, but they will typically not repeat this elsewhere.

It is the discourse model's job to accumulate the information about the identity of important problems in its knowledge representation. This can be done either via coreference or via some other mechanism that infers that the discourse is still concerned with the same problem. This may seem a very hard task, but at least it is not doomed in principle: in earlier work we managed to train non-experts in performing similar inferences and judgements during AZ annotation, using no world knowledge, only discourse cues.

5 Design of a recogniser

How could all this be recognised in unlimited text? The recognition of rhetorical moves would drive recognition with this model; as the only visible parts of the argument, rhetorical moves correspond to the bottom-up element. In contrast, high-level intentions form the top-down, *a priori* expectations. They can only ever be inferred, because the authors typically leave them implicit, so their recognition will never be made with absolute certainty.

A hybrid statistical-symbolic recogniser of scientific argumentation could instantiate the network in Fig. 2 on the fly for each new incoming paper, and keep a knowledge base of propositions derived during recognition. Whenever one of the moves is detected, the activation of its associated box is triggered. Statistically trained recognisers based on superficial features and evidence from tens of thousands of analysed papers provide a confidence value for the recognition of each move, which is translated into the strength of activation. The symbolic part of the recogniser keeps track of the logic representation accumulated up to that point in processing, and per-

forms inference as to which higher-level intention is supported by currently activated rhetorical moves.

The output of such an analysis would be a partially activated network expressing the overall argument likely to be followed in the paper, where each node in the network is annotated with a more or less instantiated knowledge representation. The activated network can be considered as an automatically-derived explanation for the place in the research space where the authors situate themselves.

Newly-derived, intermediate levels of information should be additionally available from such an analysis, as a side-effect of this hybrid style of recognition. For instance, coreference resolution is an important aspect of analysis and contributes to the superficial features. It could also feed into a mechanism that determines which of the cited previous approaches is central to the argumentation in the paper, which of these the authors present as their main rivals or collaborators, and which aspects of existing work they criticise or praise.

It is quite obvious that a solution to this task would be immediately useful for a host of applications in search, summarisation and the teaching of scientific writing. As the system would be able to associate textual statements with the corresponding likely intentions it recognised, it could produce a justification for its overall analysis of the argument. Operating as a text critiquer, such a system could point out badly-expressed instances of well-known argumentation patterns, e.g. missing or weak evidence for particular high-level intentions.

Appealing though such applications are, the main point of the analysis laid out here is the development of a theory of text understanding of naturally occurring arguments in scientific text. Given the state of current NLP technology, some of the intermediate levels of recognition necessary for this seem to us to be within reach in the near future.

6 Conclusions

This paper promotes robust text understanding of scientific articles in a deeper manner than is currently practiced, as this would lead to more informative, symbolic representations of argument structuring. Mature technologies exist for determining

specific scientific entities such as gene names (cf. the review by Campos et al., 2014) and specific events such as protein–gene interactions (e.g., Rebolz et al., 2005). In contrast to our work, such approaches are domain-specific and only recognise a small part of the entities or relationships modelled here. A different line of research associates text pieces with the research phase or information structure a given statement belongs to, where information structure is defined in terms of methods, results, conclusions etc, as in the work of Liakata et al. (2010), Guo et al. (2013) and Hirohata et al. (2008). A related task, hedge detection in science, has been established and competitively evaluated (see Farkas et al. (2010) for an overview of the respective CoNLL shared task). While these two approaches (information structuring and hedge recognition) are domain-independent like ours, the analysis presented here aims at a deeper, more informative representation of relationships between general entities in the research space.

At the other end of the spectrum, we are aware of at least one deeper analysis of argument structure in science than ours, which is manual and takes world-knowledge into account, namely Green (2014); our approach differs from hers in that we opt to model argumentation in a domain- and discipline-independent manner, which is automatic but necessarily at a far shallower level.

Our claims in this paper include that a logical scientific argument structure exists and can be interpreted by a human reader, even in light of ambiguity and although only some steps of the argumentation are explicitly stated. We have also claimed that this type of analysis holds for all disciplines in principle, but certainly for all empirical sciences. We further claim that a substantial part of the argumentation in a well-written paper is recognisable to a reader even if they do not have any domain knowledge. These are rather strong claims: It is not even clear whether humans can recognise the explicit argumentation parts, let alone the inferred ones. We therefore need to substantiate the claims with annotation experiments.

In our work to date, we have made empirical observations about argumentation structure in synthetic chemistry, computer science, computational linguistics, and agriculture, but many of these are confined to the level of AZ or CFC. We are now

in the process of corroborating the argumentation-level observations by corpus annotation of rhetorical moves. This initially takes the form of adding information to already existing AZ- and CFC-level annotation, with the aim of constructing a full-scale rhetorical move annotation. Higher-level goals will then be annotated as a second step.

Practical work also concerns building the recognisers of rhetorical moves. Several such recognisers already exist and will be refined in future work. It will be interesting to study exactly when inference about higher-level intentions becomes necessary, and which kinds of constraints can be derived from the argumentation network and the knowledge representation so as to usefully guide the inference mechanism.

References

- Charles Bazerman. 1985. Physicists reading physics, schema-laden purposes and purpose-laden schema. *Written Communication*, 2(1):3–23.
- Philippe Besnard and Anthony Hunter. 2008. *Elements of argumentation*. MIT Press.
- Christine L. Borgman and Jonathan Furner. 2002. Scholarly communication and bibliometrics. In *Annual review of information science and technology: Vol. 36*, pages 3–72. Information Today, Medford, NJ.
- Kevin W. Boyack and Richard Klavans. 2010. Co-citation analysis, bibliographic coupling, and direct citation: Which citation approach represents the research front most accurately? *Journal of the American Society for Information Science and Technology*, 61(12):2389–2404.
- Stefanie Brüninghaus and Kevin D. Ashley. 2005. Generating legal arguments and predictions from case texts. In *Proceedings of the 10th international conference on Artificial intelligence and law*, pages 65–74.
- David Campos, Srgio Matos, and Jos Lus Oliveir. 2014. Current methodologies for biomedical named entity recognition. In Mourad Elloumi and Albert Y. Zomaya, editors, *Biological Knowledge Discovery Handbook: Preprocessing, Mining and Postprocessing of Biological Data*. Wiley.
- Sandra Carberry. 1990. *Plan Recognition in Natural Language Dialogue*. MIT Press, Cambridge, MA.
- Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, editors. 1990. *Intentions in Communication*. MIT Press, Cambridge, MA.
- A.M. Cohen, W.R. Hersh, K. Peterson, and Po-Yin Yen. 2006. Reducing workload in systematic review prepa-

- ration using automated citation classification. *Journal of the American Medical Informatics Association*, 13(2):206–219.
- Robin Cohen. 1984. A computational theory of the function of clue words in argument understanding. In *Proceedings of the 10th International Conference on Computational Linguistics (COLING-84)*, pages 251–255.
- Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games. *Artificial Intelligence*, 77:321–357.
- Richrd Farkas, Veronika Vincze, Gyrgy Mra, Jnos Csirik, and Gyrgy Szarvas. 2010. The conll-2010 shared task: learning to detect hedges and their scope in natural language text. In *Proceedings of CoNLL '10: Shared Task Proceedings of the Fourteenth Conference on Computational Natural Language Learning — Shared Task*.
- Mark Garzone and Robert E. Mercer. 2000. Towards an automated citation classifier. In *Proceedings of the 13th Biennial Conference of the CSCI/SCEIO (AI-2000)*, pages 337–346.
- Nancy L. Green. 2014. Towards creation of a corpus for argumentation mining the biomedical genetics research literature. In *Proc. of the First Workshop on Argumentation Mining, ACL 2014*.
- Yufan Guo, Roi Reichart, and Anna Korhonen. 2013. Improved information structure analysis of scientific documents through discourse and lexical constraints. In *Proceedings of NAACL-2013*, Atlanta, US.
- Kenji Hirohata, Naoaki Okazaki, Sophia Ananiadou, and Mitsuru Ishizuka. 2008. Identifying sections in scientific abstracts using conditional random fields. In *Proceedings of the Third International Joint Conference on Natural Language Processing (IJCNLP 2008)*, pages 381–388, Hyderabad, India. ACL Anthology Ref. I08-1050.
- Ken Hyland. 1998. Persuasion and context: The pragmatics of academic metadiscourse. *Journal of Pragmatics*, 30(4):437–455.
- Maria Liakata, Simone Teufel, Advaith Siddharthan, and Colin Batchelor. 2010. Corpora for conceptualisation and zoning of scientific papers. In *In: Proceedings of LREC-10*, Valetta, Malta.
- Greg Myers. 1992. In this paper we report...—speech acts and scientific facts. *Journal of Pragmatics*, 17(4):295–313.
- Peter Norvig. 1989. Marker passing as a weak method for text inferencing. *Cognitive Science*, 13(4):569–620.
- Diarmuid O’Seaghdha and Simone Teufel. 2014. Unsupervised learning of rhetorical structure with untopic models. In *Proceedings of the 25th International Conference on Computational Linguistics (COLING 2014)*, Dublin, Ireland.
- Martha E. Pollack. 1986. A model of plan inference that distinguishes between the beliefs of actors and observers. In *Proceedings of the 24th Annual Meeting of the Association for Computational Linguistics (ACL-86)*, pages 207–214, New York, US.
- Martha E. Pollack. 1990. Plans as complex mental attitudes. In P.R. Cohen, J. Morgan, and M.E. Pollack, editors, *Intentions in Communication*, pages 77–103. MIT Press, Cambridge, MA.
- Dietrich Rebholz-Schuhmann, H Kirsch, and F Couto. 2005. Facts from textis text mining ready to deliver? *PLoS Biol*, 3(2). doi:10.1371/journal.pbio.0030065.
- Francoise Salager-Meyer. 1992. A text-type and move analysis study of verb tense and modality distributions in medical English abstracts. *English for Specific Purposes*, 11:93–113.
- Roger C. Schank and Robert P. Abelson. 1977. *Scripts, Goals, Plans and Understanding*. Lawrence Erlbaum, Hillsdale, NJ.
- John Swales, 1990. *Genre Analysis: English in Academic and Research Settings. Chapter 7: Research articles in English*, pages 110–176. Cambridge University Press, Cambridge, UK.
- Simone Teufel, Advaith Siddharthan, and Colin Batchelor. 2009. Towards discipline-independent argumentative zoning: Evidence from chemistry and computational linguistics. In *Proceedings of EMNLP-09*, Singapore.
- Simone Teufel. 1998. Meta-discourse markers and problem-structuring in scientific articles. In *Proceedings of the ACL-98 Workshop on Discourse Structure and Discourse Markers*, pages 43–49, Montreal, Canada.
- Simone Teufel. 2000. *Argumentative Zoning: Information Extraction from Scientific Text*. Ph.D. thesis, School of Cognitive Science, University of Edinburgh, Edinburgh, UK.
- Simone Teufel. 2010. *The Structure of Scientific Articles: Applications to Citation Indexing and Summarization*. CSLI Publications.
- Stephen Toulmin. 1958. *The Uses of Argument*. Cambridge University Press.
- Douglas Walton, Chris Reed, and Fabrizio Macagno. 2008. *Argumentation Schemes*. Cambridge University Press.
- Howard D. White. 2004. Citation analysis and discourse analysis revisited. *Applied Linguistics*, 25(1):89–116.