

Annotating Arguments in Scientific Publications

Anonymous Author

February 14, 2018

1 Introduction

Arguments are an important part of scientific writing and reflect the reasoning of researchers during the scientific process. Therefore, identifying argumentative structures in scientific publications plays a key role in understanding the scientific discourse and in making sense of massive amounts of scientific literature. It can help improving a bunch of other related computational tasks, such as summarization of scientific literature [10], research trend prediction [7] or citation analysis [6].

Hence, the goal of this annotation study is to identify argumentative components as well as their relations in scientific publications.

The present annotation guide gives an overview of the underlying argumentation theory that we are going to use and provides some examples as well as system usage instructions.

Please read the guide carefully. If you have any questions, don't hesitate to ask.

2 A Model of Argumentation

There exists many models of argumentation, such as the one introduced by Toulmin [11] or Reed and Walton [8], inter alia. What most of these models have in common is that they define several argumentative components and relationships between them [5]. To minimize the complexity of the annotations, we define the following two main components¹:

1. Claim

A *claim* corresponds to the point an arguer is trying to make. It is an assertion or a hypothesis made publicly for general acceptance and reflects therefore the author's opinion to a controversy. The claim can be identified by looking at certain key phrases such as *we think ..* or *we suggest ...*

¹Note that the examples in this section might not be completely annotated for the purpose of reducing complexity when introducing the several parts of our annotation scheme.

Paraphrasing it as *I/ We think that ..* may help to identify it.

Example: Harry is a British subject.

Interestingly, in scientific publications we can notice two types of claims: Claims that describe the scientific community or the background of the work, in contrast to claims that are made about the work of the authors themselves. To capture those differences, we introduce two subcategories:

(a) *Background_claim*

These are claims that do not relate directly to the work presented in the subject paper but that express a general believe or a general attitude regarding the domain.

Example: "The range of breathtaking realistic 3D models is only limited by the creativity of artists and resolution of devices." [1]

This sentence can be easily identified as a claim, because it corresponds to a hypothesis that the authors are putting forward and on which they are basing their arguments in the course of the article. But, this claim relates to the domain of 3D models in general and not specifically to the work the others are proposing. Therefore it has to be annotated as *background_claim*.

(b) *Own_claim*

In contrast to the sub-category *background_claim* presented before, *own_claim* is intended for capturing claims that relate directly to the work, which is the topic of the subject paper, i.e. to the authors' own work. In general, it provides more specific information.

Example: "We provide detailed reasons why and how the inverse operation can improve the results." [1]

In this example the authors are claiming that their approach is superior to others. It is therefore a claim that relates to their own work and has to be labeled as *own_claim*.

Right now we assume that *background_claims* can be typically found in the introduction of the section while *own_claims* might be more present in the middle of the paper. Usually, *background_claims* might be more difficult to identify as claims – a very conservative annotator might not annotate them at all.

2. Data

Data is the fact which we present as support for the *claim*. It is often also called evidence, ground, premise or precondition. It gives answer to the questions: What are the facts supporting the claim? What is the evidence? Why should someone believe this? In which particular case does

the claim hold?

It can for example be some kind of knowledge or an observation or even the results of an experiment. In scientific writing, citations often correspond to data, for example when the authors are referring to previous results to support their hypothesis.

Example: Harry was born in Bermuda.

Sometimes, claim and data might be difficult to distinguish and depending on the context, it might be possible that the same sentence belongs to either the one or the other category. In case you are not sure ask yourselves: Does the sentence really represent a fact? Something that is already proven (data)? Or is it rather a statement the authors are making, but that can be seen differently (claim)?

Apart from these components we define the following argumentative relationships:

1. Supports

A *Supports* relationship is a directed relationship from a component *a* to a component *b* if *a* backs *b*. Usually, this relationship exists from *data* to *claim*, but in many cases a *claim* might support another *claim*. Other combinations are still possible.

Example: "Since [data: vertex transformations can be easily implemented in the graphic card], [background_claim: SSD is very popular in circumstances that require animating a number of characters in real time]."[1]

data $\xrightarrow{\text{supports}}$ background_claim

2. Contradicts

A relationship of type *contradicts* represents the counterpart of the *supports* relationship, but in contrast, it is a bi-directional, i.e., symmetric relationship. An instance of this type of relationship exists between a component *a* and a component *b*, if component *a* contradicts *b* and vice versa.

Example: "Given physical principles, [background_claim 1: this category can generate more believable animation effects compared to its geometric counterpart]. But [background_claim 2: they are seldom applied to interactive applications] because of the high cost of computing and complicated algorithms." [1]

background_claim 1: $\xleftrightarrow{\text{contradicts}}$ background_claim 2:

In general, the skeleton of a very simple argument can be illustrated graphically as shown in figure 1 or using natural language as follows:

D, so C.

Fitting in the example dealing with Harry, we get the following argument:

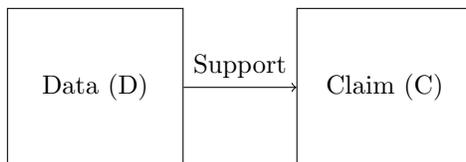


Figure 1: A simple argument structure.

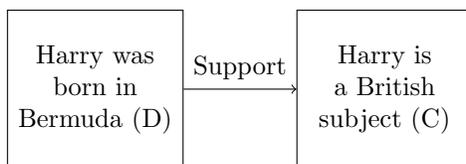


Figure 2: The simple argument structure applied to an example.

Harry was born in Bermuda, so Harry is a British subject.

An illustration of this argument structure applied to the example can be seen in figure 2. In the simplest case, the relationships presented above exist between a data component and a claim. But there might be also several data components connected to a claim and a more complex argument might also include relationships between claims. A more compound argument is illustrated in figure 3. In case you have identified multiple claims inside the same publication, the claims themselves may also have a certain relationship with each other. For example, if you think for that the claims support each other (imagine a major claim and a minor claim), you should connect the claims with a *supports* relation, such that an inter-argument structure gets visible.

In addition to the two argumentative inter-component relationships presented above, we have to specify a third "artificial" type of relation, which exist only inside a single component, *parts_of_same*, as well as a fourth non-argumentative relationship, *semantically_same*.

- *Parts_of_same*

In real-world examples of argumentative structures a single component, such as for example a claim, might be split up in several parts. Nevertheless, we would like to be able to recognize that those parts actually belong to the same component and should be treated in that way. We specify a new type of relationship, being bidirectional, intra-component, and non-argumentative for identifying such discontinuous components.

Example: [*background_claim 1: Tense interpretation has received much attention in linguistics*] (Partee (1984), Hinrichs (1986), Nerbonne (1986), inter alia), [*background_claim 2: and natural language processing*] (Weber (1988), Kameyama et al.(1993), Lascarides and Asher (1993), inter

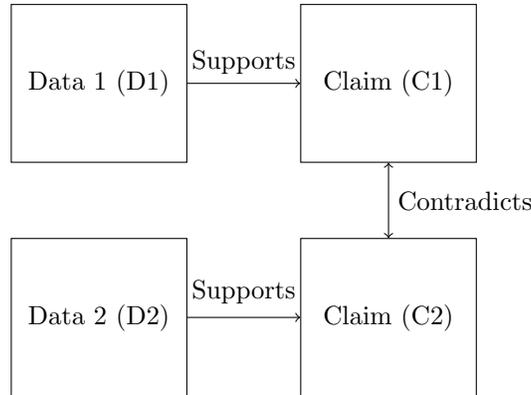


Figure 3: A complex argument: Claim C2 is contradicting claim c1; both claims are supported by corresponding data.

alia).” [4]

background_claim 1 $\xleftrightarrow{\text{parts_of_same}}$ background_claim 2

- Semantically_same

When presenting work in science, the authors often repeat certain parts of their argumentation to underline their importance or to remind the reader of what was said before with regard to what is coming next. This means, that we could find multiple instances of a single argumentative component, e.g. maybe the concrete wording of two text spans might differ, but they are essentially expressing the same claim. To capture those cases, we introduce a symmetric and non-argumentative relationship, *semantically_same*.

Example: ”This section will describe the implementation of our inverse algorithm and why [*own_claim 1*: it is an improvement]. (...) We call the PSD scheme without the inverse operation as “forward PSD”, and comparison to it will be used to demonstrate [*own_claim 2*: the superiority of the inverse method].” [1]

own_claim 1 $\xleftrightarrow{\text{semantically_same}}$ own_claim 2

In case you encounter more than two instances of essentially the same component you do not have to draw the relationship to all instances, but to the last instance of that component that you have already connected.

To sum up, the component types which you can choose during the annotation process are *background_claim*, *own_claim*, and *data*. In addition, you may assign two types of argumentative relationships, *supports* and *contradicts*, as well as two non-argumentative relationships, *parts_of_same* and *semantically_same*.

△ Note that we do not want to annotate whole sentences only, i.e. we would

like to study annotation on a fine-grained level in which single argumentative components might consist of a single clause or even of a citation marker only. Therefore, there might be also several components of different types present in the same sentence. Please always try to annotate the minimal text span and omit conjunctions, such as "because", when they are not part of an argumentative component but rather connecting multiple. The annotated span must be understandable on its own.

Example: "But they are seldom applied to interactive applications because of the high cost of computing and complicated algorithms." [1]
Annotate as claim "they are seldom applied to interactive applications" but not "seldom applied to interactive applications" (also no "but" or "because" in the span).

3 The Annotation Process

When performing the annotation, we recommend the following process:

1. Read the whole publication from the beginning to the end to get an overview about the structure of the paper as well as about the overall content.
2. Then, start from the introduction section, i.e., you can leave out the abstract², and read sentence by sentence carefully.
3. For each sentence³:
 - (a) Decide: Does it contain argumentative components or not?
 - (b) If yes, which parts of the sentence, i.e., text spans, of the sentence are argumentative (the whole sentence, a relative clause, etc.)?
 - (c) Whenever you have identified an argumentative part(s), decide whether they are of type *data* or *claim* and in case you identify a claim, choose one of the two subcategories *background_claim* and *own_claim*.
 - (d) Next, in case you have identified several parts: Do they actually belong together but are just split because of linguistic reasons? If yes, assign a *parts_of_same* relationship.
 - (e) For each argumentative component identified think about other possible relationships that might hold between other already identified components.
 - (f) If you identify such a relationship: Of which type is it (*supports*, *contradicts* or *semantically_same*)?
4. Finally, go again through your annotations and check whether they make sense.

²Similarly, you can leave out the appendix and the acknowledgements.

³Keep in mind that in theory, components might also consist of more than one sentence!

△ If you encounter problems or difficult cases, please collect them, such that we can discuss them in our training sessions. Please copy the corresponding text spans and paste them together with your questions or thoughts in our *discussion file*.

4 Consistency Rules and other Remarks

In order to increase the consistency of the annotations, please stick to the following rules when annotating.

1. Punctuation

Whenever you mark a text span and this span ends with a punctuation (comma, period, semicolon, colon, etc.) do not include this symbol into your annotation.

Example: "Therefore, [*own_claim*: providing a flexible and efficient solution to animation remains an open problem]." [1]

2. References

As said before, references are often used as data. Whenever you encounter this constellation, make sure that you only mark the minimal span and omit brackets, parentheses etc.

Example: "[*background_claim*: A nice review of SSD is given] in [*data*: 1]]." [1]

3. Preconditions

We would like to capture preconditions for claims as data. Cue phrases indicating preconditions are for example "if", "for", "in the case of", "whenever", "when" etc. But naturally there might be constellations, in which these phrases are not related to a precondition. In case they are, make sure that you only mark the precondition itself and not the cue phrase.

Example: "For [*data*: small deformations, [*own_claim*: both algorithms produce similar results], as in the [*data*: second row of Figure 6] (...)." [3]

4. Examples

Sometimes, the authors use examples to support their argumentation. Therefore, in many cases, examples correspond to data. Similar to preconditions, we only want to capture the examples and not the cue phrases, such as "such as", "e.g.", "for example", "for instance" etc. Also, please mark each example separately.

Example: "Each time when [*data*: a frame goes wrong], [*background_claim*: a production cannot afford major revisions] such as [*data*: resculpting models] or [*data*: re-rigging skeletons]." [1]

5. Figures and Tables

Especially when presenting results, the authors often refer to tables or figures in order to support their claims. In such a case, a figure corresponds to data. Again, make sure that you only mark the minimal spans.

Note that this also applies to other specific parts of the publications, e.g., sections.

Example: "As explained in the [*data: next section*], [*own_claim: the SBS works on a circular arc instead of segment*], see [*data: Figure 1*]."

6. Structure-related hints

The overall structure of a scientific publication might already give you a hint which argumentative components you might encounter. Of course, you should not rely on this and you might encounter many variations. Nevertheless, it might help to keep in mind which part of the paper you are currently annotating. The structure of a computer science paper and the types of components you will find are typically like this:

- Introduction: Mostly *background_claims*, but in the end *own_claims*
- Related Works: Mostly *background_claims*, often references as *data*
- Method or Implementation: In the beginning maybe some *background_claims*, but mostly *own_claims*; it might have less argumentative components, as authors are often relatively descriptive in this part
- Experiments or Results: Mostly *own_claims*, often figures and tables as *data*
- Conclusion: Some *background_claims*, mostly *own_claims*.

5 Examples

5.1 Example 1

Example taken from [2]:

"Although [*background_claim: the scheme was originally proposed as a preferred way of implementing the independence assumption*], we suggest that [*own_claim 1: it is also appropriate for implementing similarity-based models, as well as class-based models*]. [*own_claim 2: It enables us to rely on direct maximum likelihood estimates*] when [*data 1: reliable statistics are available, and only [data 2: otherwise]*] [*own_claim 3: resort to the estimates of an 'indirect' model*]."

Relations:

background_claim $\xleftrightarrow{\text{contradicts}}$ own_claim 1
own_claim 2 $\xrightarrow{\text{supports}}$ own_claim 1
data 1 $\xrightarrow{\text{supports}}$ own_claim 2
data 2 $\xrightarrow{\text{supports}}$ own_claim 3
own_claim 3 $\xrightarrow{\text{supports}}$ own_claim 1

In this example the phrase "we suggest" indicates that the authors are trying to make an argumentative point. We can easily reformulate this part of the example as: We think that it [the model] is also appropriate for .. . Also, this statement is an important part of the authors' work. Therefore, it should be labeled as *own_claim*. But it *contradicts* a claim that was made about the related works (*background_claim*) before. Furthermore, after making their statement, the authors explain why they think so, i.e. they provide evidence for convincing the reader of their statement. Therefore, that sentence has to be labeled as *data*. In addition to annotating both components, please indicate the relationship between them by drawing a relation of type *supports* from the data to the claim.

5.2 Example 2

Example taken from [4]:

[*background_claim 1: Tense interpretation has received much attention in linguistics*] ([*data 1: Partee (1984)*], [*data 2: Hinrichs (1986)*], [*data 3: Nerbonne (1986)*], inter alia), [*background_claim 2: and natural language processing*] ([*data 4: Webber (1988)*], [*data 5: Kameyama et al. (1993)*], [*data 6: Lascarides and Asher (1993)*], inter alia).

Relations:

background_Claim 1 $\xrightarrow{\text{parts_of_same}}$ background_claim 2
 data 1 $\xrightarrow{\text{supports}}$ background_Claim 1
 data 2 $\xrightarrow{\text{supports}}$ background_Claim 1
 data 3 $\xrightarrow{\text{supports}}$ background_Claim 1
 data 4 $\xrightarrow{\text{supports}}$ background_claim 2
 data 5 $\xrightarrow{\text{supports}}$ background_claim 2
 data 6 $\xrightarrow{\text{supports}}$ background_claim 2

Interestingly, in this example, the authors do not announce that a claim is made. There is no clear linguistic clue. Furthermore, regarding the discourse structure of the publication, the phrase "Tense interpretation has received much attention in linguistics [...] and natural language processing [...]" is not the central hypothesis. Instead, it's just some kind of background information the authors are adding, such that the readers understand that the topic is important. Nevertheless, when looking at the micro-structure of arguments, we can find argumentative components here and as we can easily rephrase the sentence as "We think that tense interpretation has received much attention in linguistics .. " we should label the phrase therefore as *background_claim*.

Another interesting aspect of this example is, that proof for the claim is not provided in terms of numbers or long textual explanations but just in form of references to other publications that are provided as examples. Therefore, the

references have to be labeled as *data*. More specifically, as each citation provides already some kind of evidence on its own, each reference has to be marked as a single component. Note that we also have to add a relation which indicates that the two parts of the claim actually belong together. Therefore, you should draw a relation of type *parts_of_same* from *background.Claim 1* to *background.Claim 2*. Note that we only want to annotate the references and not the parentheses around them.

5.3 Example 3

Example taken from [4]:
 ”We argue that [*own_claim*: aspects of both analyses are necessary to account for the recovery of temporal relations.] To demonstrate our approach we will address the following examples; passages [...] are taken from Lascarides and Asher (1993).”

Here, the claim is clearly announced by the phrase ”We argue that .. ”. The authors state that they are now going to make their major claim, which they gonna argue for using some examples. But one has to be careful: The authors do not explain the facts why they think their claim holds, instead they just announce that they are going to do so. Therefore, the second part of the excerpt must not be labeled as *data*.

5.4 Example 4

Example taken from [4]:
 ”[*own_claim 1*: The tenses used may not completely specify the implicit temporal relations between the described events]. We claim that [*own_claim 2*: these relations may be further refined by constraints imposed by the coherence relation operative between clauses]. We describe three coherence relations relevant to the examples in this paper and give temporal constraints for them.”

Relations:

own_claim 1 $\xrightarrow{\text{supports}}$ *own_claim 2*

Here, we encounter two claims regarding the method of the authors. Therefore, both have to be marked as *own_claim*. The second claim is announced by ”We claim that [...]” but the first one is more difficult to identify. Still, the first claim provides reason why the relations should be further refined. Therefore, a relationship of type *supports* needs to be drawn.

5.5 Example 5

Example taken from [1]:

”For the situation where other unknown skinning operations are adopted, we propose a unified framework which will be discussed in the following section.”

In this example, nothing should be annotated because the authors are just saying that they propose something, but they are not stating any particular property about the framework.

5.6 Example 6

Example taken from [3]:

”[own_claim 1: The proposed skin deformation system is by no means perfect]; [own_claim 2: it cannot compete with complex, layered models]. However, [own_claim 3: the SBS algorithm offers reasonable price for elimination of the notorious LBS artifacts.]”

Relations:

$$\begin{array}{l} \text{own_claim 2} \xrightarrow{\text{supports}} \text{own_claim 1} \\ \text{own_claim 1} \xleftarrow{\text{contradicts}} \text{own_claim 3} \\ \text{own_claim 2} \xleftarrow{\text{contradicts}} \text{own_claim 3} \end{array}$$

Here, the authors are making three points: First of all, they claim that their system is not perfect. As this directly relates to their own work, this statement has to be annotated as *own_claim*. Similarly, the second and the third claim are of type *own_claim*. Moreover, the second claim provides reason for the first claim. Therefore, *own_claim 2* supports *own_claim 1*. In contrast, the third claim is introduced with the adverb ”however”, which corresponds to a hint that the following statement is inconsistent with what was said before. Hence, *own_claim 3* contradicts the two former claims.

5.7 Example 7

Example taken from [3]:

”First, if [*data*: we substitute r c in place of v], [*own_claim 1*: no rotation occurs], which means that [*own_claim 2*: r c is] indeed [*own_claim 3*: a center of rotation].”

Relations:

$$\begin{array}{l} \text{data} \xrightarrow{\text{supports}} \text{own_claim 1} \\ \text{own_claim 1} \xrightarrow{\text{supports}} \text{own_claim 2} \\ \text{own_claim 2} \xleftrightarrow{\text{parts_of_same}} \text{own_claim 3} \end{array}$$

This sentence exposes a reasoning structure, which is indicated by ”if .. ”. In such cases, we would like to capture the precondition introduced by ”if” as *data* (note that we do not want to annotate the ”if” itself), while the second part, which in this case is not explicitly marked by ”then”, has to be annotated as *claim*. Here, the authors relate to their own work, which is why it is of type *own_claim*. As the part starting with ”if” represents a precondition, it supports the second part. Furthermore, the authors follow from these two parts that ”[.] r c is indeed a center of rotation.” The word ”indeed” just corresponds to a figure of speech. Therefore, we exclude it from the text span marking only the rest as claims and connecting those two latter parts with *parts_of_same*. The first claim supports the whole construct involving *own_claim 2* and *own_claim 3*, but as those two are essentially one claim, we only draw the supports relation to the part that’s closest from *own_claim 1*, *own_claim 2*.

5.8 Example 8

Example taken from [3]:

”Let us denote matrices by capital letters, while vectors and quaternions by bold.”

This sentence is non-argumentative as the authors are just explaining a mathematical convention in order to make the reader better understand the formulas. Nothing has to be annotated.

5.9 Example 9

Example taken from [1]:

”We reformulate the problem as $f(y_i) = v_i + w_i - \text{SKINNING}_i(v_r + d_i)2 + \lambda w_i^2$, where y_i is a concatenated vector $y_i = [d_i, w_i]$ and λ is an arbitrary small number.”

Similar to the example above, nothing should be annotated here as the authors are just describing their method.

5.10 Example 10

Example taken from [1]:

”Basic skinning provided by Maya is called in the loop of minimization scheme.”

Again, nothing should be annotated here as the authors are just describing their implementation.

5.11 Example 11

Example taken from [1]:

”*[background_claim: Skeleton Subspace Deformation (SSD) is the predominant approach to character skinning at present]*.”

In contrast to the examples above, the authors are not just describing a method in this sentence, they are making an assessment of it and basically stating that everybody in their specific community is using it. They are presenting it as a fact, but people could argue about this and the authors could provide supporting evidence for their opinion. Therefore, this has to be annotated as *background_claim*.

6 Using the Annotation Tool: BRAT

In the present annotation study we decided for using the BRAT Rapid Annotation Tool (BRAT) [9].

When launching BRAT you will first see a welcome message, which links you to user manual of the system ⁴ and provides basic information regarding the usage of BRAT. You may find it helpful to read the usage instructions before starting to annotate. Nevertheless, in case you encounter any problems with the system, please contact us.

⁴<http://brat.nlplab.org/manual.html>

The following sections provide a brief overview of the main features you will need.

6.1 Selecting a Document

After you have closed the welcome message the "Open" dialog appears, in which you can select the document you would like to annotate. You can also open this dialog at any time by selecting the button "Collection", which appears together with the rest of the main menu, when you hover the top area of the screen.

In that dialog you can now select the folder, in which you would like to work. *Please only open the folder, which we assigned to you beforehand!*

After opening the folder, you can select one of the documents, which you would like to annotate. You can also go back to documents you already worked on and change your annotations if you want to. We recommend that you keep track of your individual annotation status, such that you have a clear overview on which publications still require your annotation.

In case you have already opened a document, you can move to the succeeding/preceding document by using the right/ left arrows in the left corner of the screen.

6.2 Logging in

You can check all the documents including the annotations at any time, but you can only modify your labels while being signed in to the system. To do so, you have to hover over the top bar of the screen such that the main menu appears and click on the "Login" button in the right corner. Then the "Login" dialog appears, which prompts you for the user name and the password, which we assigned to you.⁵ You have to submit your credentials by pressing the button "ok".

6.3 Annotating a Document

After you have logged in to the system and selected a document, you can start to annotate!

BRAT distinguished between two types of annotations that you can make: *entities* and their *associations*. In our case, entities correspond to argumentative components, e.g. *own_claim* or *data*, while associations correspond to argumentative relations such as *supports*. You can stop/ pause the annotation process at any point in time as your progress will be saved. Nevertheless, we recommend to always finish the annotation of a publication before moving ahead or taking a longer break.

⁵In case you forget your credentials, don't hesitate to contact us.

6.3.1 Annotating Entities

To annotate an argumentative component, e.g. an entity, you should first select (mark) the text span which corresponds in your opinion to the argumentative component. Next, a dialog will ask you for the type of argumentative component you just selected ("entity type"). To do so, just mark the radio button to the left of the type with which you would like to label the text span. Furthermore, you can optionally leave some comments on the annotation. To confirm your annotation, close the dialog by pressing the button "ok", otherwise press "cancel".

6.3.2 Annotating Associations

To annotate an association in BRAT, you first must have annotated the argumentative components, i.e. entities, between which you would like to draw the relation. Then, select the entity *from* which the relation starts and hold the mouse down while moving to the entity *to* which the relation goes. The movement corresponds to a drag-and-drop action, which you may know from other applications.

Next, the association dialog opens, in which you can specify the type of relation similar to specifying the type of entity before. You can confirm your annotation by pressing "ok", otherwise press "cancel".

6.3.3 Editing Annotations

To edit your annotations, double-click on the highlighted area. The "Edit Annotation" dialog will open, which displays several options, such as changing the type of an annotation or deleting it.

6.3.4 Troubleshooting

If you encounter any problems with or weird behavior of the system, just contact us in order to find a solution. Sometimes, also the BRAT manual might give you a hint to the solution. Known problems are the following:

- *Problem:* The file can be seen, but annotations can not be made. E.g., after marking a span, the window to select the label for the annotation does not pop up.
Solution: Log in.
- *Problem:* After marking a span, the window always jumps to the beginning of the page.
Solution: Reload the page.
- *Problem:* BRAT indicates that it cannot visualize an annotation because of inconsistencies.
Solution: Contact us, such that we can clean up your annotation file.

- *Problem:* BRAT does not show anything (maybe just a weird line) when I open a file.
Solution: Contact us, such that we can clean up your annotation file.

7 Concluding Remarks

Identifying argumentative structures in scientific publications is an important step in the direction of making sense of massive amounts of scientific literature and gaining a deeper understanding of the sociology of science.

We would like to thank all the annotators for their efforts.

References

- [1] A powell optimization approach for example-based skinning in a production animation environment. *Computer animation and social agents*, pages 141–150, 2006.
- [2] Ido Dagan, Fernando Pereira, and Lillian Lee. Similarity-based estimation of word cooccurrence probabilities. In *Proceedings of the 32Nd Annual Meeting on Association for Computational Linguistics*, ACL '94, pages 272–278, Stroudsburg, PA, USA, 1994. Association for Computational Linguistics.
- [3] Ladislav Kavan and Jiří Žára. Spherical blend skinning: A real-time deformation of articulated models. In *Proceedings of the 2005 Symposium on Interactive 3D Graphics and Games*, I3D '05, pages 9–16, New York, NY, USA, 2005. ACM.
- [4] Andrew Kehler. Temporal relations: reference or discourse coherence? In *Proceedings of the 32nd annual meeting on Association for Computational Linguistics*, pages 319–321. Association for Computational Linguistics, 1994.
- [5] Christian Kirschner, Judith Ecker-Köhler, and Iryna Gurevych. Linking the thoughts: Analysis of argumentation structures in scientific publications. In *ArgMining@ HLT-NAACL*, pages 1–11, 2015.
- [6] Anne Lauscher, Goran Glavaš, Simone Paolo Ponzetto, and Kai Eckert. Investigating convolutional networks and domain-specific embeddings for semantic classification of citations. In *Proceedings of WOSP 2017*, volume tba, page tba, Toronto, 2017. ACM ; Mannheim Univ. Press.
- [7] Kathy McKeown, Hal Daume, Snigdha Chaturvedi, John Paparrizos, Kapil Thadani, Pablo Barrio, Or Biran, Suvarna Bothe, Michael Collins, Kenneth R. Fleischmann, Luis Gravano, Rahul Jha, Ben King, Kevin McInerney, Taesun Moon, Arvind Neelakantan, Diarmuid O’Seaghdha, Dragomir

- Radev, Clay Templeton, and Simone Teufel. Predicting the Impact of Scientific Concepts Using Full-text Features. *J. Assoc. Inf. Sci. Technol.*, 67(11):2684–2696, November 2016.
- [8] Chris Reed and Douglas Walton. Argumentation schemes in argument-as-process and argument-as-product, 2003.
- [9] Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii. Brat: A web-based tool for nlp-assisted text annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, EACL '12, pages 102–107, Stroudsburg, PA, USA, 2012. Association for Computational Linguistics.
- [10] Simone Teufel and Marc Moens. Summarizing scientific articles: Experiments with relevance and rhetorical status. *Comput. Linguist.*, 28(4):409–445, December 2002.
- [11] Stephen E. Toulmin. *The Uses of Argument, Updated Edition*. Cambridge University Press, 2008.