

# Automated Statement Extraction from Press Briefings

Jüri Keller,<sup>1</sup> Meik Bittkowski,<sup>2</sup> Philipp Schaer<sup>3</sup>

**Abstract:** Scientific press briefings are a valuable information source. They consist of alternating expert speeches, questions from the audience and their answers. Therefore, they can contribute to scientific and fact-based media coverage. Even though press briefings are highly informative, extracting statements relevant to individual journalistic tasks is challenging and time-consuming. To support this task, an automated statement extraction system is proposed. Claims are used as the main feature to identify statements in press briefing transcripts. The statement extraction task is formulated as a four-step procedure. First, the press briefings are split into sentences and passages, then claim sentences are identified through sequence classification. Subsequently, topics are detected, and the sentences are filtered to improve the coherence and assess the length of the statements.

The results indicate that claim detection can be used to identify statements in press briefings. While many statements can be extracted automatically with this system, they are not always as coherent as needed to be understood without context and may need further review by knowledgeable persons.

**Keywords:** Computational Journalism; Claim Detection; Data Mining; Natural Language Processing

## 1 Introduction

Scientific press briefings are a valuable instrument in scientific communication. They consist of alternating expert speeches and answers to questions from the audience, which let scientists and journalists immediately and jointly address the information needs of journalists. The SMC press briefings<sup>4</sup> used in this work, typically begin with an introduction of the participating experts followed by moderated questions from journalists and answers from invited experts. Additionally, the key results are concluded in closing statements. Press briefings can directly contribute to a more fact-based media coverage by connecting journalists and scientists. Even though press briefings are highly informative, filtering this information and extracting relevant statements remains challenging and time-consuming due to the high entropy and domain-specific language used.

In this context, the research question will be answered: To what extent can claim detection be used to extract statements from scientific press briefings?

---

<sup>1</sup> Technische Hochschule Köln, Claudiusstraße 1, 50678 Köln, Germany jueri.keller@smail.th-koeln.de

<sup>2</sup> Science Media Center Germany, Schloss-Wolfsbrunnenweg 33, 69118 Heidelberg, Germany bittkowski@sciencemediacenter.de

<sup>3</sup> Technische Hochschule Köln, Claudiusstraße 1, 50678 Köln, Germany philipp.schaer@th-koeln.de

<sup>4</sup> [https://www.sciencemediacenter.de/alle-angebote/alle-angebote/?tx\\_solr%5Bfilter%5D%5B0%5D=type%3APress+Briefing](https://www.sciencemediacenter.de/alle-angebote/alle-angebote/?tx_solr%5Bfilter%5D%5B0%5D=type%3APress+Briefing)

To approach this domain-specific problem and automatically identify relevant statements from press briefings, a pipeline is proposed, relying on claims as a central element of a statement. A transformer based language model [CSM20] is used to classify claim sentences, which are subsequently filtered and clustered to improve the relevance and coherence of the emerging statements.

For example, a sentence like *"We do not yet know what a long covid course of Omicron infection looks like."* is considered a *complete claim*, while the sentence *"If I only have mild symptoms, it doesn't mean I won't have problems in the long run."* is only considered an *incomplete claim* as it can not be understood without context.

The main contribution of this work is the application of state-of-the-art methods in claim detection to develop a system that extracts statements from transcripts of scientific press briefings. Additionally, a novel dataset of 53 transcribed German press briefings is created and partially annotated. The dataset and implementation are made publicly available via GitHub<sup>56</sup> and Zenodo [Ke22].

## 2 Related Work

Automated extraction of statements from press briefings occurs at the intersection of journalism and computer science. While the premise and context emerge from journalism, the technical methods originate from computer science.

Various definitions of a claim exist [Da17; LEC16]. Generically, a claim indicates the author's position related to a main concept (topic) of a sentence [Le17]. Therefore, claims are indicated as the main feature of a statement. Generically, a claim is a short phrase people can have different opinions on [LEC16; LT16]. Since claims are conceptually complex, it is often difficult to determine what a claim is and what not [LT15]. To clarify, this work refers to a claim as a sentence asserting something on the main concept.

Claim detection is an information extraction task often part of a data processing pipeline solving more complex problems. Automatically detecting claims in texts has a variety of applications in fields like decision-making, argument mining, fact-checking, or document processing [Da17; Le18]. As diverse as the applications for claim detection are, so are the domains it is applied in, such as for example political, debate, legal, and the web [LT16].

Besides binary classification, different claim types or the part of the sentence that is a claim can be identified [Ko18; LT16]. Levy et al. [Le14] first formally defined the task of automated claim detection for computational argumentation while working on automatic argument mining [Le14; SI21]. The systems developed in the following years can be categorized by their ability to detect claims without further context or the scope of application, e.g., a

---

<sup>5</sup> Press Briefing Claim Dataset at GitHub: [https://github.com/jueri/press\\_briefing\\_claim\\_dataset](https://github.com/jueri/press_briefing_claim_dataset)

<sup>6</sup> Statement Extractor at GitHub: [https://github.com/jueri/statement\\_extractor](https://github.com/jueri/statement_extractor)

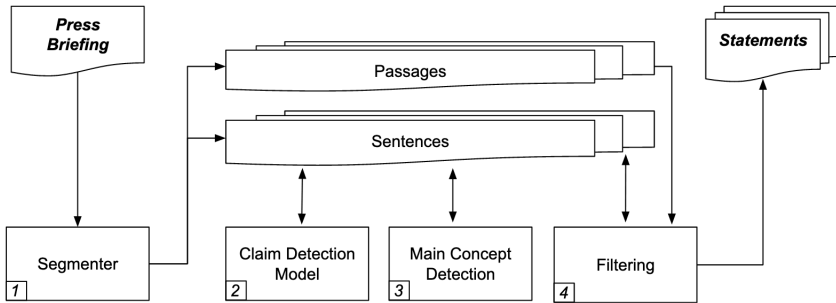


Fig. 1: Schematically visualization of the pipeline approach.

corpus of documents or a single one [Da17; LT15]. Chakrabarty et al. and Daxenberger et al. provide more extensive overviews of different claim detection models and datasets [CHM19; Da17].

### 3 Statement Detection System

Figure 1 visualizes the pipeline approach to extract statements from press briefings. The press briefing needs to be initially segmented into passages (1). Then, at the system’s core, a classification model is used to identify sentences containing claims (2). Further, the main concept of the press briefing and the sentences are detected (3). This information is used in the last step to filter out the sentences that are not or incomplete claims (4).

**1. Segmentation:** The press briefings are split by a sentence tokenizer, and neighboring and coherent sentences are combined into statements of multiple sentences. Therefore, the sentences are represented in vector space through spaCy [HMB20] and neighboring sentences are clustered by similarity with the TextSplit [AG15] algorithm. Starting with a list of vector representations of sentences, TextSplit calculates a score for each segment. By modifying the segment boundaries, the score changes. The optimal segmentation is determined by minimizing the aggregated segment scores [AG15; Mo20].

**2. Claim Detection:** A German language model [CSM20] based on the BERT [De19] architecture is fine-tuned for the task of sentence classification. The model’s hyperparameters are tuned to achieve the best results on available training data. The German BERT version GBERT is used as the foundational model for the classifier [CSM20]. This model is fine-tuned as a claim sentence classifier using a dataset of 3000 sentences created for this purpose from the publicly available SMC press briefings<sup>7</sup>. The dataset contains a total of 53 press briefings from a time span of over four years. It consists of 24,897 sentences with an average length of 17.31 tokens.

<sup>7</sup> <https://www.sciencemediacenter.de/alle-angebote>

**3. Main Concept Detection:** Based on the intuitive assertion that a claim sentence with a topic similar to the overall topic of the press briefing the sentence originates from is most relevant, the topics of both the sentences and the overall press briefings (i.e., the title and the introduction text) are detected. Therefore, Wikipedia articles are used as representations for the topic of a sentence or the whole press briefing. All Wikipedia articles related to tokens in the sentences and the title and the introduction texts are detected. The sentence is considered more relevant if a sentence can be linked to the same Wikipedia articles as the title or introduction text. For this task, the wikification APIs Dandelion<sup>8</sup> and TagMe<sup>9</sup> are used [FS10].

**4. Sentence Filtering:** Two techniques are applied to filter out claim sentences with a low topical similarity between the main concept of the sentence and the overall main concept to improve the coherence and relevance. The similarity is measured by two approaches based on embeddings or Wikipedia concepts. Both approaches create a similarity score that can be used for filtering with a minimum threshold.

The first approach calculates similarity by creating a vector representation from the title of the press briefing and the individual sentences and then measuring the cosine similarity between both. The vector representations are created by combining spaCy [HMB20] word embeddings. The second approach is based on the Wikipedia concepts detected as described previously. The confidence scores of shared Wikipedia concepts between the title or introduction text and the individual sentence are summed to create a score of topical relatedness.

## 4 Experimental Evaluation

Three experiments were conducted to evaluate the system's ability to extract statements from press briefings. Therefore, three new press briefings with 799 sentences from different categories were annotated as ground truth.

**Claim Detection:** For hyperparameter optimization, models with different configurations were trained for six epochs. The optimal number of epochs was determined in conjunction with the learning rate by comparing the evaluation loss rates after each epoch. Learning rates, ranging from 0.00005 to 0.01, were tested. The best results could be achieved with a learning rate of  $1e-5$  and three epochs. For the final system, a model with an F1 Score of 0.89, a precision of 0.92 and a recall of 0.86 was used. This model assigned slightly more false positive than false negative labels, but since the human in the loop can decide on the misclassified sentences, this behavior is preferred.

To evaluate the claim detection component of the system, claim sentences classified with minimal confidence scores of 0.7, 0.8 and 0.9 are analyzed independently. Since the

---

<sup>8</sup> <https://dandelion.eu/>

<sup>9</sup> <https://sobigdata.d4science.org/web/tagme/tagme-help>

Tab. 1: Results for detected *complete claims*. The first three columns hold the results for different confidence thresholds of the claim detection model. The last three columns contain the results of a system using a claim detection model threshold of 0.8 and the different main concept methods. The best results are highlighted in **bold**.

Confidence	0.9	0.8	0.7	0.8 embedding	0.8 w. title	0.8 w. intro
F1	0.466	<b>0.473</b>	0.450	<b>0.481</b>	0.339	0.341
Precision	<b>0.426</b>	0.378	0.339	<b>0.463</b>	0.456	0.430
Recall	0.513	0.632	<b>0.671</b>	<b>0.500</b>	0.270	0.283

confidence score limits the number of sentences considered as claims by the model, a score small enough to allow claim sentences to be classified and high enough to exclude non-claim sentences leads to the best results. The overall best results for detecting claims independent of their completeness were achieved at a confidence of 0.7 with an F1 score of 0.68, a precision of 0.89 and a recall of 0.55. By investigating individual results for the claim types, *complete claims* (that can stand on their own) and *incomplete claims* (that can not), it can be assessed if the filtering and sentence clustering methods improve the coherence. In total, a maximum of 167 of 224 *incomplete claims* and 102 of 152 *complete claims* claim sentences could be identified with various system configurations. With the claim detection model and minimum confidence of 0.8, the highest F1 score of 0.47 could be achieved. Raising the confidence threshold to 0.9 results in a higher precision of 0.43. Decreasing the score to 0.7 leads to a higher recall of 0.67.

**Statement Filtering:** Based on the results of the first experiments, the claim sentences with a confidence score of 0.8 are further filtered as this threshold achieved the best F1 score. To exclude incoherent statements, the similarity between the sentence’s main concept and the press briefing’s main concept is calculated. The main concepts of the press briefing based on *embeddings*, *title wikification* and *introduction wikification* are evaluated individually. For each method, a minimum similarity is chosen by investigating the similarity distribution for bends. With the *introduction wikification* based filtering, a precision of 0.43 was achieved. The *title wikification* method reaches a precision of 0.456, and the *embedding* method exceeds this score with the best precision of 0.463. All results are presented in table 1.

**Sentence Clustering:** The last experiment evaluates the sentence clustering method. Claim sentences with a confidence of at least 0.9 are enlarged to statements of multiple sentences. The statements are then assessed according to their coherence. To measure if the larger statements created by the similarity-based sentence clustering are more coherent, the *complete claim* and *incomplete claim* ratio for all methods are compared. With 53 % *complete claims*, the coherence of the statements extracted exceeds all other methods and the baseline. 46 % of the statements extracted with the claim detection module and a minimum confidence score of 0.9 are complete statements. The embedding method of the main concept module extracts 50 % *complete claims*.

## 5 Discussion and Conclusion

The analysis shows that the press briefings are very argumentative content full of claims. While most claims are incomplete and can only be understood with further information, 46 % can stand alone. The sentences extracted by the proposed system are highly likely to be claims, although not necessarily complete claims. This finding is supported by the low false-positive rate of the system considering all claims. Missing coherence is the main error of the statements extracted.

Compared to the initial results of the claim detection model with an accuracy of 0.89, the results achieved in the experiments are noticeably lower. This may be caused by the dataset used for training, which only contains *complete claims* and *no claims* and should be improved in future work. By adding a filter based on main concepts, some incoherent claim sentences could be successfully excluded, and the precision improved. A smaller confidence threshold of the initial claim detection module could improve these results further by providing more claim sentences in general. The embedding-based main concept method provided the best results. The Wikification-based methods come to an extent as some topics can only be expressed poorly by Wikipedia concepts, and other Wikipedia concepts may not exist or be captured by the model.

Assessing statements of multiple sentences, some statements previously classified as *incomplete claims* can be understood without additional information through the added neighboring sentences. By systematically adding sentences, more coherent statements could be created. However, this method sacrifices the conciseness of statements since they get longer.

In summary, this work investigates how statements can be mined from scientific press briefings. A dataset is created from publicly available press briefings and used to fine-tune a claim detection model. Additional methods are tested to improve the quality of the extracted statements. More context is added to the statements, and the statements are filtered based on their topic. The results are compared to a gold standard to evaluate the system's performance.

The results show that the system can differentiate between claims of any type and *no claim* sentences with high precision (0.89), but only half of the claims can be recalled (0.55). Considering only *complete claim* sentences, the system's overall performance decreases (F1 0.47), but relatively more complete claims can be identified. These results emphasize the system's difficulty differentiating between complete and incomplete claim sentences. The resulting statements gain coherence by adding more context. Similarly, incomplete claim sentences can be filtered out by low topical similarity, which improves the results slightly. The evaluations show that the presence of a claim can be used as an indicator to detect statements. However, the results leave room for improvement in many aspects. Especially the coherence of the extracted statements needs to be improved to extract statements that can stand on their own.

## References

- [AG15] Alemi, A. A.; Ginsparg, P.: Text segmentation based on semantic word embeddings. CoRR abs/1503.05543/, 2015, URL: <http://arxiv.org/abs/1503.05543>.
- [CHM19] Chakrabarty, T.; Hidey, C.; McKeown, K.: IMHO fine-tuning improves claim detection. In (Burstein, J.; Doran, C.; Solorio, T., eds.): Proceedings of the 2019 conference of the north american chapter of the association for computational linguistics: Human language technologies, NAACL-HLT 2019, minneapolis, MN, USA, june 2-7, 2019, volume 1 (long and short papers). Association for Computational Linguistics, pp. 558–563, 2019, URL: <https://doi.org/10.18653/v1/n19-1054>.
- [CSM20] Chan, B.; Schweter, S.; Möller, T.: German’s next language model. In (Scott, D.; Bel, N.; Zong, C., eds.): Proceedings of the 28th international conference on computational linguistics, COLING 2020, barcelona, spain (online), december 8-13, 2020. International Committee on Computational Linguistics, pp. 6788–6796, 2020, URL: <https://doi.org/10.18653/v1/2020.coling-main.598>.
- [Da17] Daxenberger, J.; Eger, S.; Habernal, I.; Stab, C.; Gurevych, I.: What is the essence of a claim? Cross-domain claim identification. In (Palmer, M.; Hwa, R.; Riedel, S., eds.): Proceedings of the 2017 conference on empirical methods in natural language processing, EMNLP 2017, copenhagen, denmark, september 9-11, 2017. Association for Computational Linguistics, pp. 2055–2066, 2017, URL: <https://doi.org/10.18653/v1/d17-1218>.
- [De19] Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In (Burstein, J.; Doran, C.; Solorio, T., eds.): Proceedings of the 2019 conference of the north american chapter of the association for computational linguistics: Human language technologies, NAACL-HLT 2019, minneapolis, MN, USA, june 2-7, 2019, volume 1 (long and short papers). Association for Computational Linguistics, pp. 4171–4186, 2019, URL: <https://doi.org/10.18653/v1/n19-1423>.
- [FS10] Ferragina, P.; Scaiella, U.: TAGME: on-the-fly annotation of short text fragments (by wikipedia entities). In: Proceedings of the 19th ACM international conference on Information and knowledge management - CIKM '10. ACM Press, Toronto, ON, Canada, p. 1625, 2010, ISBN: 978-1-4503-0099-5, URL: <http://portal.acm.org/citation.cfm?doid=1871437.1871689>, visited on: 10/07/2021.
- [HMB20] Honnibal, M.; Montani, I.; Boyd, S. V. L. A.: spaCy: Industrial-strength Natural Language Processing in Python, version 0.1, Feb. 2020, URL: <https://doi.org/10.5281/zenodo.1212303>.

- [Ke22] Keller, J.: Automated Statement Extraction from Press Briefings, version 0.1, Feb. 2022, URL: <https://doi.org/10.5281/zenodo.6047551>.
- [Ko18] Konstantinovskiy, L.; Price, O.; Babakar, M.; Zubiaga, A.: Towards automated factchecking: Developing an annotation schema and benchmark for consistent automated claim detection. CoRR abs/1809.08193/, 2018, URL: <http://arxiv.org/abs/1809.08193>.
- [Le14] Levy, R.; Bilu, Y.; Hershcovich, D.; Aharoni, E.; Slonim, N.: Context dependent claim detection. In (Hajic, J.; Tsujii, J., eds.): COLING 2014, 25th international conference on computational linguistics, proceedings of the conference: Technical papers, august 23-29, 2014, dublin, ireland. ACL, pp. 1489–1500, 2014, URL: <https://aclanthology.org/C14-1141/>.
- [Le17] Levy, R.; Gretz, S.; Sznajder, B.; Hummel, S.; Aharonov, R.; Slonim, N.: Unsupervised corpus-wide claim detection. In (Habernal, I.; Gurevych, I.; Ashley, K. D.; Cardie, C.; Green, N.; Litman, D. J.; Petasis, G.; Reed, C.; Slonim, N.; Walker, V. R., eds.): Proceedings of the 4th workshop on argument mining, ArgMining@EMNLP 2017, copenhagen, denmark, september 8, 2017. Association for Computational Linguistics, pp. 79–84, 2017, URL: <https://doi.org/10.18653/v1/w17-5110>.
- [Le18] Levy, R.; Bogin, B.; Gretz, S.; Aharonov, R.; Slonim, N.: Towards an argumentative content search engine using weak supervision. In (Bender, E. M.; Derczynski, L.; Isabelle, P., eds.): Proceedings of the 27th international conference on computational linguistics, COLING 2018, santa fe, new mexico, USA, august 20-26, 2018. Association for Computational Linguistics, pp. 2066–2081, 2018, URL: <https://aclanthology.org/C18-1176/>.
- [LEC16] Liebeck, M.; Esau, K.; Conrad, S.: What to do with an airport? Mining arguments in the german online participation project tempelhofer feld. In: Proceedings of the third workshop on argument mining, hosted by the 54th annual meeting of the association for computational linguistics, ArgMining@ACL 2016, august 12, berlin, germany. The Association for Computer Linguistics, 2016, URL: <https://doi.org/10.18653/v1/w16-2817>.
- [LT15] Lippi, M.; Torrioni, P.: Context-independent claim detection for argument mining. In (Yang, Q.; Wooldridge, M. J., eds.): Proceedings of the twenty-fourth international joint conference on artificial intelligence, IJCAI 2015, buenos aires, argentina, july 25-31, 2015. AAAI Press, pp. 185–191, 2015, URL: <http://ijcai.org/Abstract/15/033>.
- [LT16] Lippi, M.; Torrioni, P.: Argumentation mining: State of the art and emerging trends. ACM Trans. Internet Techn. 16/2, 10:1–10:25, 2016, URL: <https://doi.org/10.1145/2850417>.
- [Mo20] Mody, N.: Finding the best part of your podcast to promote via NLP, en, June 2020, URL: <https://towardsdatascience.com/finding-the-best-part-of-your-podcast-to-promote-via-nlp-f844a88b287a>, visited on: 09/08/2021.



- [SI21] Slonim, N.; Bilu, Y.; Alzate, C.; Bar-Haim, R.; Bogin, B.; Bonin, F.; Choshen, L.; Cohen-Karlik, E.; Dankin, L.; Edelstein, L.; Ein-Dor, L.; Friedman-Melamed, R.; Gavron, A.; Gera, A.; Gleize, M.; Gretz, S.; Gutfreund, D.; Halfon, A.; Hershovich, D.; Hoory, R.; Hou, Y.; Hummel, S.; Jacovi, M.; Jochim, C.; Kantor, Y.; Katz, Y.; Konopnicki, D.; Kons, Z.; Kotlerman, L.; Krieger, D.; Lahav, D.; Lavee, T.; Levy, R.; Liberman, N.; Mass, Y.; Menczel, A.; Mirkin, S.; Moshkovich, G.; Ofek-Koifman, S.; Orbach, M.; Rabinovich, E.; Rinott, R.; Shechtman, S.; Sheinwald, D.; Shnarch, E.; Shnayderman, I.; Soffer, A.; Spector, A.; Sznajder, B.; Toledo, A.; Toledo-Ronen, O.; Venezian, E.; Aharonov, R.: An autonomous debating system. en, *Nature* 591/7850, pp. 379–384, Mar. 2021, ISSN: 0028-0836, 1476-4687, URL: <http://www.nature.com/articles/s41586-021-03215-w>, visited on: 10/20/2021.