

# Factuality Detection using Machine Translation - a Use Case for German Clinical Text

Mohammed Bin Sumait, Aleksandra Gabryszak, Leonhard Hennig and Roland Roller

German Research Center for Artificial Intelligence (DFKI)

Speech and Language Technology Lab

{firstname.lastname}@dfki.de

## Abstract

Factuality can play an important role when automatically processing clinical text, as it makes a difference if particular symptoms are explicitly not present, possibly present, not mentioned, or affirmed. In most cases, a sufficient number of examples is necessary to handle such phenomena in a supervised machine learning setting. However, as clinical text might contain sensitive information, data cannot be easily shared. In the context of factuality detection, this work presents a simple solution using machine translation to translate English data to German to train a transformer-based factuality detection model.

## 1 Introduction

Factuality refers to the concept that a speaker can present statements about world events with varying degrees of uncertainty as to whether they happened. Factuality reflects, for instance, if an event is affirmed, negated, or uncertain. In the medical domain, detecting if symptoms or diseases are signaled as present, not present, possibly or doubtfully present, and therefore uncertain is essential. Detecting factuality is challenging since it can be expressed by very different linguistic categories (e.g. verbs, nouns, adjectives, adverbs), plus it must be taken into account how they are embedded in a sentence (Rudinger et al., 2018a). Additionally, linguistic factuality cues can be very domain-specific, so the availability of relevant datasets is essential.

Classical supervised machine learning requires training data, and, at the same time, most existing datasets are published in English. In addition, clinical text contains sensitive patient data, which often makes it difficult to share due to ethical and legal aspects. Although the situation has slowly changed regarding the availability of German clinical text resources (Modersohn et al., 2022), many other languages suffer a similar situation. Conversely, the quality of machine translation has significantly improved in the last decade, also regarding the trans-

lation of biomedical text/publications, including clinical case reports (Neves et al., 2022). For this reason, this work explores the usage of machine translation to create (translated) text resources for factuality detection in German clinical text.

Clinical notes are short text documents written by physicians during or shortly after the treatment of a patient. In general, this kind of text contains much valuable information about the current health condition, as well as treatment, of the patient. They differ from biomedical publications and clinical case reports, as notes are often written under time pressure with a high information density, a telegraphic writing style, non-standardized abbreviations, colloquial errors, and misspellings. Therefore, it is unclear if current machine translation systems can handle this text, considering that data might contain sensitive information and should not be shared with a third party outside the hospital.

This work makes the following contributions: 1) We successfully use a local machine translation to train a model for factuality detection on German clinical text. 2) Our model outperforms the only ‘competitor’ NegEx, and 3) will be published as open access model<sup>1</sup>. Finally, 4) for those interested in NegEx, we release it as a modular PyPI package with a few important fixes<sup>2</sup> and also propose improvement suggestions to the used trigger sets.

## 2 Methods and Data

The idea of this work is based on the usage of machine translation to generate a German corpus to train a classifier dealing with factuality in clinical text. In the following, we outline the approach, the necessary methods, and the dataset used.

### 2.1 Factuality Detection

In literature, (medical) factuality detection is often reduced to a simple classification. Given a sentence

<sup>1</sup><https://huggingface.co/binsumait/factual-med-bert-de>

<sup>2</sup><https://github.com/DFKI-NLP/pynegex>

Factuality	English	German translation
affirmed	Clinically, a <E>severe neuropsychological syndrome</E> was found when the patient was taken over.	Klinisch fand sich bei Übernahme des Patienten in <E>schweres neuropsychologisches Syndrom</E>.
negation	Patient denies <E>headache</E>.	Patient verneint <E>Kopfschmerzen</E>.
possible	Thus, a <E>tumour</E> cannot be ruled out.	Ein <E>Tumor</E> kann daher nicht ausgeschlossen werden.

Table 1: Example sentences with target entities, factuality label, and possible translations.

and an entity, the task is to define the factuality of the entity in the given context. In most cases, the entity of interest is a symptom or medical condition. Most related work targets the three classes **affirmed**, **negated** and **possible**. However, as simple as this sounds, factuality cannot always be easily mapped to those few classes.

One of the most prominent tools to deal with factuality in the medical text is NegEx (Chapman et al., 2001), a rule-based approach with pre-defined regular expressions, so-called triggers, and can detect the three aforementioned factuality classes. It achieves, particularly in the context of negations, quite good results on clinical text. Hedges instead offer more possibilities for how they are described, therefore achieving a much lower performance. Initially, it was developed for English, but over the years, it has also been translated into other languages, such as Spanish or Swedish (Cotik et al., 2016b; Chapman et al., 2013). In addition, many alternative (machine learning) solutions have been published in the last two decades. We refer to the overview by Khandelwal and Sawant (2019) for more details. For German, however, only one negation detection exists, which relies on the NegEx solution and uses a set of translated trigger words (English to German) (Cotik et al., 2016a).

## 2.2 Data

In the following, we briefly introduce the data used for this work. First, we present i2b2, which has been used for machine translation and to train our model. In addition, we later test our model on additional German data, namely Ex4CDS and NegEx-Ger, and in the appendix also BRONCO150.

The 2010 i2b2/VA data (Uzuner et al., 2011) consists of English medical text and includes three tasks - extraction of concepts, assertions identification, and relation detection. In this work, we focus on the assertion task. Overall a total of six assertion types were considered, namely present, absent, possible, conditional, hypothetical and not associated with the patient. However, this work focused only on the first three labels, as only those are considered within NegEx. i2b2 data is translated to

German to train a German machine learning model.

**Ex4CDS** (Roller et al., 2022) is a small dataset of physicians’ notes containing explanations in the context of clinical decision support. The notes are written in German and include various annotation layers, including factuality. As the data includes multiple factuality labels, we reduced the labels to our three target labels, mapping *possible-future* and *unlikely* to *possible*, and *minor* to *affirmed*. As target entities, we consider only sentences containing *medical-conditions*.

**NegEx-Ger** is a small dataset consisting of sentences taken from clinical notes and discharge summaries and has been used initially to evaluate the German NegEx version in Cotik et al. (2016a). For our use case, the data has been used for testing, and for this, we merged the sentences of both clinical text types. However, the number of sentences containing the possible label is small (22 for discharge summaries and 4 for clinical notes).

## 2.3 Translation Approach

For our proposed idea, two aspects need to be considered: First, we aim at a solution that could be applied to sensitive data. Therefore, the machine translation component must run locally. This means we cannot rely on the variety of existing state-of-the-art online approaches. Second, as we define factuality as a classification problem with a given sentence (context) and an entity, our translations need to keep track of the target entity within a sentence. A simple example is given in Table 1, which shows an English sentence with a target entity ‘headache’ and the label ‘negation’. The German translation needs to keep the focus on the target entity.

In this work, we rely on TransIns (Steffen and van Genabith, 2021), an open-source machine translation that can be installed locally. TransIns is built on MarianNMT (Junczys-Dowmunt et al., 2018) framework and enables translating texts with an embedded markup language. Specifically, we translate sentences with tagged entities, as shown in Table 1.

A manual inspection revealed multiple problems

with the translations: In some cases (roughly 40% of the issues), translations were corrupt as they contained cryptic and/or repetitive text sequences that were foreign from the original text. Such noise patterns could partially or entirely affect the target texts’ context. Or, in very few cases (only 4%), no translation output could be produced. In the rest of the cases, the markup no longer included the target entity. In any way, such output has been discarded from the data, and we resulted in 18,297 data points (initially 18,397), which we used to train and evaluate our machine learning model.

### 3 Experiments and Results

We conduct three different experiments - starting with the English i2b2 data, we use Bio+Discharge Summary BERT (Alsentzer et al., 2019) and compare the results to NegEx. Similar experiments have also been conducted in other papers. However, in our case, those results serve as a comparison. Thus, the model is not optimized to achieve the best possible performance. Next, we train German-MedBERT (Shrestha, 2021) on the translated i2b2 data and compare the results to the performance of the German NegEx implementation. Finally, we apply both German factuality approaches to different German medical texts to determine how well the models perform in a more realistic setup.

Label	NegEx			BERT-based		
	Prec	Rec	F1	Prec	Rec	F1
E Affirmed	0.88	0.97	0.93	<b>0.97</b>	<b>0.99</b>	<b>0.98</b>
N Negated	0.89	0.79	0.84	<b>0.98</b>	<b>0.97</b>	<b>0.97</b>
G Possible	0.79	0.04	0.08	<b>0.85</b>	<b>0.64</b>	<b>0.73</b>
G Affirmed	0.84	0.96	0.90	<b>0.96</b>	<b>0.98</b>	<b>0.97</b>
E Negated	0.83	0.65	0.73	<b>0.95</b>	<b>0.93</b>	<b>0.94</b>
R Possible	0.28	0.02	0.04	<b>0.80</b>	<b>0.64</b>	<b>0.71</b>

Table 2: Performance results between NegEx baselines and BERT-based models on the original English i2b2 dataset (upper part) and German translation (lower part).

The results of the first two experiments are presented in Table 2 and show various interesting findings: Firstly, NegEx provides impressive results on the affirmed label, good results for negations, and unsatisfying results for the possible label. Moreover, on both datasets, English and German, the BERT-based model outperforms NegEx, on all scores. Additionally, results on the English dataset are always higher than those on the translated dataset. This might be unsurprising as data quality decreases. Finally, the table shows that BERT-based models show a substantial increase in

performance for the possible label.

Label	NegEx			BERT-based		
	Prec	Rec	F1	Prec	Rec	F1
N Affirmed	0.96	0.94	0.95	<b>0.97</b>	<b>0.96</b>	<b>0.96</b>
E Negated	0.93	0.96	0.95	<b>0.97</b>	<b>0.98</b>	<b>0.97</b>
G Possible	0.46	0.50	0.48	<b>0.50</b>	0.50	<b>0.50</b>
E Affirmed	0.85	0.88	0.86	<b>0.88</b>	<b>0.92</b>	<b>0.90</b>
X Negated	0.66	0.89	0.76	<b>0.86</b>	<b>0.95</b>	<b>0.90</b>
4 Possible	0.50	0.18	0.26	<b>0.61</b>	<b>0.38</b>	<b>0.47</b>

Table 3: Performance results on different German medical text sources, namely the original German NegEx (upper part), and Ex4CDS dataset (lower part).

Table 3 presents the performance of the NegEx and the BERT-based model on two German datasets. In the upper part of the table, the results on NegEx-Ger are presented and the results on Ex4CDS are in the lower part. Similarly, as on the translated i2b2 dataset in Table 2, the machine learning model outperforms NegEx. However, this time the performance gain is not so strong anymore. The NegEx-Ger is small and relatively homogeneous (regarding the variety of negations), and NegEx already performs well on the negations. Therefore the machine learning model achieves only a performance boost of two points in F1. In case of possible, the number of examples might be too small to see the benefit of the ML model.

On Ex4CDS data, NegEx already struggles with *negated* (0.76) and performs low in the case of *possible* (0.26) - although the results are much better in comparison to the results on i2b2 (English and German). Here, the machine learning model leads to a performance boost of 14 points for *negated* and 21 points for *possible*.

### 4 Analysis and Discussion

Our results indicate that we can successfully apply machine translation to generate a German clinical dataset to train a machine learning model with. Most notably, this model can outperform NegEx, which partially already provides satisfying results. While it is important that a negation detection tool for German clinical text needs to run within a hospital infrastructure, it might be questionable if BERT-based approaches might be the right solution, as it requires much more hardware resources than the simple NegEx solution. This is supported by the results on NegEx-Ger, in which the BERT achieves only a minor performance gain. However, as this data is small and homogeneous, the results on Ex4CDS affirm the usage of machine learning,

as we achieve a notable performance gain. Note, information about the frequency of each label in the test data is provided in the appendix. As our BERT model was trained on potential suboptimal translations, we analyse some errors in more detail in the following.

#### 4.1 Linguistic Error Analysis

Our analysis focuses on the prediction errors caused by the translation or by differences in the features of the German and English language. Table 7 contains full-text examples illustrating the issues described below.

In various cases, a factuality cue was completely missing in the translation, or the sense of the cue was not preserved (e.g., *to rule out* was translated with *Vorschriften* instead of *ausschließen*). In those cases, NegEx and BERT labeled the instances wrongly as affirmations.

In other cases, we observe that the factuality cues are outside of the original data’s entities but in the translation they are placed within the entity markup. That is often correlated with the prediction changing from negation or possible to affirmation. For example, both NegEx and BERT correctly recognized the negated assertion of the original phrase *did not notice [any blood]*, whereas both German models consider the translation *bemerkte [kein Blut]* as affirmed in which the negation cue (*not / kein*) became part of the entity.

For NegEx, a further problem are missing factuality cues in the trigger list. For example, it systematically does not recognize the cue *verleugnen* (one of the possible translations of the word *deny*, which is included in the English NegEx). Additionally, some problems with factuality cues are specific to the German language and require additional handling: (a) German compounds must be written as one word; unfortunately, German NegEx cannot handle cases when a compound consists of words referring to a medical problem and its negation (e.g. *schmerzfrei / pain free*), since it seems not to recognize a factuality cue if it is not written as a separate phrase, (b) cues with umlauts in text such as *aufgelöst* seem not to be recognized, because the umlauts are encoded as *oe* in the German trigger list, (c) missing possible word orders of factuality phrases (e.g. word order might depend on the embedding syntactic structure; e.g. *wurde ausgeschlossen* vs. *ausgeschlossen wurde* in a main vs. subordinate clause).

## 5 Related Work

**Machine Translation for Cross-lingual Learning** MT is a popular approach to address the lack of data in cross-lingual learning (Hu et al., 2020; Yarmohammadi et al., 2021). There are two basic options - translating target language data to a well-resourced source language at inference time and applying a model trained in the source language (Asai et al., 2018; Cui et al., 2019), or translating source language training data to the target language, while also projecting any annotations required for training, and then training a model in the target language (Khalil et al., 2019; Kolluru et al., 2022; Frei and Kramer, 2023). Both approaches depend on the quality of the MT system, with translated data potentially suffering from translation or alignment errors (Aminian et al., 2017; Ozaki et al., 2021). While the quality of machine translation for health-related texts has significantly improved (Neves et al., 2022), using MT in the clinical domain remains underexplored, with very few exceptions (Frei and Kramer, 2023).

**Factuality Detection** Previous research focused mainly on assigning factuality values to events and often framed this task as a multiclass classification problem over a fixed set of uncertainty categories (Rudinger et al., 2018b; Zerva, 2019; Pouran Ben Veyseh et al., 2019; Qian et al., 2019; Bijl de Vroe et al., 2021; Vasilakes et al., 2022). In the biomedical/clinical domain, Uzuner et al. (2011) present the i2b2 dataset for assertion classification, and Thompson et al. (2011) introduce the Genia-MK corpus, where biomedical relations have been annotated with uncertainty values. van Aken et al. (2021) release factuality annotation of 5000 data points sourced from MIMIC. Kilicoglu et al. (2017) introduce a dataset of PubMed abstracts with seven factuality values, and find that a rule-based model is more effective than a supervised machine learning model on this dataset.

## 6 Conclusion

This work presented a machine learning-based factuality detection for German clinical text. The model was trained on translated i2b2 data and tested, first on the translations and then on other German datasets and outperformed an existing method for German, NegEx. The simple machine translation approach might interest the Non-English clinical text processing community. The model will be made publicly available.

## Ethical Considerations

We use the original datasets “as is”. Our translations of i2b2 thus reflect any biases of the original dataset and its construction process, as well as biases of the MT models (e.g., rendering gender-neutral English nouns to gendered nouns in German). We use BERT-based PLMs in our experiments, which were pretrained on a large variety of medical source data. Our models may have inherited biases from these pretraining corpora.

Since medical data is highly sensitive with respect to patient-related information, all datasets used in our work are anonymized. The authors of the original datasets (Uzuner et al., 2011; Roller et al., 2022) have stated various measures that prevent collecting sensitive, patient-related data. Therefore, we rule out the possible risk of sensitive content in the data.

## Limitations

A key limitation of this work is the dependence on a machine translation system to get high-quality translations and annotation projections of the source language dataset. Depending on the availability of language resources and the quality of the MT model, the translations we use for training and evaluation may be inaccurate, or be affected by translation noise, possibly leading to overly optimistic estimates of model performance. In addition, since the annotation projection is completely automatic, any alignment errors of the MT system will yield inaccurate instances in the target language.

## Acknowledgements

This research was supported by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) through the project KEEPHA (442445488) and the German Federal Ministry of Education and Research (BMBF) through the projects KIBATIN (16SV9040) and CORA4NLP (01IW20010).

## References

- Emily Alsentzer, John R Murphy, Willie Boag, Weihung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical bert embeddings](#). *arXiv preprint arXiv:1904.03323*.
- Maryam Aminian, Mohammad Sadegh Rasooli, and Mona Diab. 2017. [Transferring semantic roles using translation and syntactic information](#). In *Proceedings*

*of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 13–19, Taipei, Taiwan. Asian Federation of Natural Language Processing.

- Akari Asai, Akiko Eriguchi, Kazuma Hashimoto, and Yoshimasa Tsuruoka. 2018. [Multilingual extractive reading comprehension by runtime machine translation](#). *ArXiv*, abs/1809.03275.
- Sander Bijl de Vroe, Liane Guillou, Miloš Stanojević, Nick McKenna, and Mark Steedman. 2021. [Modality and negation in event extraction](#). In *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*, pages 31–42, Online. Association for Computational Linguistics.
- Wendy W Chapman, Will Bridewell, Paul Hanbury, Gregory F Cooper, and Bruce G Buchanan. 2001. [A simple algorithm for identifying negated findings and diseases in discharge summaries](#). *Journal of biomedical informatics*, 34(5):301–310.
- Wendy W Chapman, Dieter Hilert, Sumithra Velupillai, Maria Kvist, Maria Skeppstedt, Brian E Chapman, Michael Conway, Melissa Tharp, Danielle L Mowery, and Louise Deleger. 2013. Extending the negex lexicon for multiple languages. *Studies in health technology and informatics*, 192:677.
- Viviana Cotik, Roland Roller, Feiyu Xu, Hans Uszkoreit, Klemens Budde, and Danilo Schmidt. 2016a. Negation detection in clinical reports written in german. In *Proceedings of the Fifth Workshop on Building and Evaluating Resources for Biomedical Text Mining (BioTxtM2016)*, pages 115–124.
- Viviana Cotik, Vanesa Stricker, Jorge Vivaldi, and Horacio Rodríguez Hontoria. 2016b. Syntactic methods for negation detection in radiology reports in spanish. In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing, BioNLP 2016: Berlin, Germany, August 12, 2016*, pages 156–165. Association for Computational Linguistics.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2019. [Cross-lingual machine reading comprehension](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1586–1595, Hong Kong, China. Association for Computational Linguistics.
- Johann Frei and Frank Kramer. 2023. [German medical named entity recognition model and data set creation using machine translation and word alignment: Algorithm development and validation](#). *JMIR Form Res*, 7:e39077.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation](#). In *Proceedings of the 37th International*

- Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. [Marian: Fast neural machine translation in C++](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Talaat Khalil, Kornel Kiełczewski, Georgios Christos Chouliaras, Amina Keldibek, and Maarten Versteegh. 2019. [Cross-lingual intent classification in a low resource industrial setting](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6419–6424, Hong Kong, China. Association for Computational Linguistics.
- Aditya Khandelwal and Suraj Sawant. 2019. NegBERT: a transfer learning approach for negation detection and scope resolution. *arXiv preprint arXiv:1911.04211*.
- Halil Kilicoglu, Graciela Rosemblat, and Thomas C. Rindfleisch. 2017. Assigning factuality values to semantic relations extracted from biomedical research literature. *PLoS ONE*, 12.
- Madeleine Kittner, Mario Lamping, Damian T Rieke, Julian Götze, Bariya Bajwa, Ivan Jelas, Gina Rüter, Hanjo Hautow, Mario Sängler, Maryam Habibi, et al. 2021. Annotation and initial evaluation of a large annotated german oncological corpus. *JAMIA open*, 4(2):oab025.
- Keshav Kolluru, Muqeeth Mohammed, Shubham Mittal, Soumen Chakrabarti, and Mausam . 2022. [Alignment-augmented consistent translation for multilingual open information extraction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2502–2517, Dublin, Ireland. Association for Computational Linguistics.
- Luise Modersohn, Stefan Schulz, Christina Lohr, and Udo Hahn. 2022. [Grascco - the first publicly shareable, multiply-alienated german clinical text corpus](#). *Studies in health technology and informatics*, 296:66–72.
- Mariana Neves, Antonio Jimeno Yepes, Amy Siu, Roland Roller, Philippe Thomas, Maika Vicente Navarro, Lana Yeganova, Dina Wiemann, Giorgio Maria Di Nunzio, Federica Vezzani, Christel Gerardin, Rachel Bawden, Darryl Johan Estrada, Salvador Lima-lopez, Eulalia Farre-maduel, Martin Krallinger, Cristian Grozea, and Aurelie Neveol. 2022. [Findings of the WMT 2022 biomedical translation shared task: Monolingual clinical case reports](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 694–723, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Hiroaki Ozaki, Gaku Morio, Terufumi Morishita, and Toshinori Miyoshi. 2021. [Project-then-transfer: Effective two-stage cross-lingual transfer for semantic dependency parsing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2586–2594, Online. Association for Computational Linguistics.
- Amir Pouran Ben Veyseh, Thien Huu Nguyen, and Dejing Dou. 2019. [Graph based neural networks for event factuality prediction using syntactic and semantic structures](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4393–4399, Florence, Italy. Association for Computational Linguistics.
- Zhong Qian, Peifeng Li, Qiaoming Zhu, and Guodong Zhou. 2019. [Document-level event factuality identification via adversarial neural network](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2799–2809, Minneapolis, Minnesota. Association for Computational Linguistics.
- Roland Roller, Aljoscha Burchardt, Nils Feldhus, Laura Seiffe, Klemens Budde, Simon Ronicke, and Bilgin Osmanodja. 2022. An annotated corpus of textual explanations for clinical decision support. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2317–2326.
- Rachel Rudinger, Aaron Steven White, and Benjamin Van Durme. 2018a. [Neural models of factuality](#). *CoRR*, abs/1804.02472.
- Rachel Rudinger, Aaron Steven White, and Benjamin Van Durme. 2018b. [Neural models of factuality](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 731–744, New Orleans, Louisiana. Association for Computational Linguistics.
- Manjil Shrestha. 2021. Development of a language model for medical domain. master thesis, Hochschule Rhein-Waal.
- Jörg Steffen and Josef van Genabith. 2021. [TransIns: Document translation with markup reinsertion](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 28–34. Association for Computational Linguistics.
- Paul Thompson, Raheel Nawaz, John McNaught, and Sophia Ananiadou. 2011. Enriching a biomedical event corpus with meta-knowledge annotation. *BMC Bioinformatics*, 12:393 – 393.

Özlem Uzuner, Brett R South, Shuying Shen, and Scott L DuVall. 2011. 2010 i2b2/va challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556.

Betty van Aken, Ivana Trajanovska, A. Siu, M. Mayrdorfer, Klemens Budde, and Alexander Loeser. 2021. Assertion detection in clinical notes: Medical language models to the rescue? *Proceedings of the Second Workshop on Natural Language Processing for Medical Conversations*.

Jake Vasilakes, Chrysoula Zerva, Makoto Miwa, and Sophia Ananiadou. 2022. [Learning disentangled representations of negation and uncertainty](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8380–8397, Dublin, Ireland. Association for Computational Linguistics.

Mahsa Yarmohammadi, Shijie Wu, Marc Marone, Haoran Xu, Seth Ebner, Guanghui Qin, Yunmo Chen, Jialiang Guo, Craig Harman, Kenton Murray, Aaron Steven White, Mark Dredze, and Benjamin Van Durme. 2021. [Everything is all it takes: A multi-pronged strategy for zero-shot cross-lingual information extraction](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1950–1967, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Chrysoula Zerva. 2019. *Automatic Identification of Textual Uncertainty*. Ph.D. thesis, University of Manchester.

## A Appendix

The main contribution of this short paper was to show that it is possible to develop a machine learning-based factuality detection for non-English, without training examples in the given language - just by using a local machine translation. In addition, we would like to present a small ‘bonus’ experiment, which did not fit into the main article anymore. More precisely, we wanted to find out how the performance of such a model changes if data in a reasonable size is available for training. The additional experiment is presented in Appendix A.1, followed by some additional text examples for the linguistic error analysis and some further information.

### A.1 Additional Experiment

The additional experiment has been conducted with the **BRONCO150** (Kittner et al., 2021) dataset, a relatively large corpus originating from 150 German oncological de-identified discharge summaries

and annotated for multiple tasks, including factuality detection. For our experiment, we consider only the target entities *diagnosis*. Similar to Ex4CDS, it has various factuality values, which we mapped to our three target labels, namely *possible future* and *speculation* to *possible*. Note, BRONCO150 contains various fragmented entities (entities split into two to three parts). For our experimental setup, we merged entity fragments and considered only those sentences with not more than 50 characters between the fragments.

The label distribution of the obtained BRONCO150 data and the distribution of the other datasets from the main paper are presented in Table 5.

First, we run the same experiment as presented in Table 3, also on BRONCO150 data. The results using our FactualMedBERT-DE model are presented in Table 4.

Label	NegEx			BERT-based		
	Prec	Rec	F1	Prec	Rec	F1
N Affirmed	0.96	0.94	0.95	<b>0.97</b>	<b>0.96</b>	<b>0.96</b>
E Negated	0.93	0.96	0.95	<b>0.97</b>	<b>0.98</b>	<b>0.97</b>
G Possible	0.46	0.50	0.48	<b>0.50</b>	0.50	<b>0.50</b>
E Affirmed	0.85	0.88	0.86	<b>0.88</b>	<b>0.92</b>	<b>0.90</b>
X Negated	0.66	0.89	0.76	<b>0.86</b>	<b>0.95</b>	<b>0.90</b>
4 Possible	0.50	0.18	0.26	<b>0.61</b>	<b>0.38</b>	<b>0.47</b>
B Affirmed	0.87	0.96	0.91	<b>0.88</b>	<b>0.97</b>	<b>0.92</b>
R Negated	0.69	0.66	0.68	<b>0.75</b>	<b>0.80</b>	<b>0.77</b>
O Possible	0.68	0.24	0.36	<b>0.73</b>	<b>0.25</b>	<b>0.37</b>

Table 4: Performance results on different German medical text sources, namely the original German NegEx (upper part), the Ex4CDS dataset (middle) and BRONCO150 (lower part).

	Affirmed	Negated	Possible
2010 i2b2/va	7603	2305	595
Ex4CDS	892	225	179
NegEx-Ger	645	443	26
BRONCO150	3179	331	523

Table 5: Support numbers in the evaluation sets for each processed dataset.

Next, we train two additional models, one on a BRONCO150 training split and a second using the BRONCO150 train together with the translated i2b2 data. Both models were initialized from the same model as that of FactualMedBERT-DE. Table 6 compares our FactualMedBERT-DE against the other two BERT-based models on the different datasets.

**Brief discussion:** The results show that each model performs best on the data of the same dataset

Model	Label	2010 i2b2/VA			NegEx-Ger			Ex4CDS			BRONCO150		
		Prec	Rec	F1	Prec	Rec	F1	Prec	Rec	F1	Prec	Rec	F1
FactualMedBERT-DE	Affirmed	0.96	0.98	<b>0.97</b>	0.97	0.96	0.96	0.88	0.92	0.90	0.88	0.97	0.92
	Negated	0.95	0.93	<b>0.94</b>	0.97	0.98	0.97	0.86	0.95	0.90	0.76	0.79	0.78
	Possible	0.80	0.64	<b>0.71</b>	0.50	0.50	0.50	0.61	0.38	0.47	0.68	0.19	0.30
BRONCO150-BERT	Affirmed	0.88	0.95	0.92	0.97	0.92	0.94	0.90	0.90	0.90	0.96	0.96	0.96
	Negated	0.95	0.67	0.79	0.97	0.97	0.97	0.89	0.88	0.88	0.95	0.83	<b>0.89</b>
	Possible	0.42	0.47	0.44	0.28	0.65	0.39	0.56	0.59	0.58	0.76	0.84	<b>0.80</b>
i2b2+BRONCO150 BERT	Affirmed	0.94	0.98	0.96	0.98	0.95	0.96	0.90	0.94	<b>0.92</b>	0.95	0.98	<b>0.97</b>
	Negated	0.96	0.91	0.93	0.98	0.97	0.97	0.90	0.91	<b>0.91</b>	0.93	0.83	0.88
	Possible	0.82	0.54	0.65	0.39	0.73	<b>0.51</b>	0.70	0.54	<b>0.61</b>	0.85	0.74	0.79

Table 6: Performance results of three BERT models trained on translated i2b2 (FactualMedBERT-DE), BRONCO150 and 2010 i2b2 + BRONCO150, respectively. The models were evaluated on different German medical text sources, namely our translated i2b2 2010 test set, the German NegEx, the Ex4CDS dataset and BRONCO150 test set. For each dataset, best per-label F1-performances are displayed in **bold**.

Issue	English	German
missing trigger in translation	The patient radiated down her left arm associated with some nausea, <u>no</u> shortness of breath, cough, vomiting, diarrhea.	Die Patientin strahlte in Verbindung mit Übelkeit, Atemnot, Husten, Erbrechen, Durchfall nach unten.
incorrect trigger translation	<u>RULE OUT FOR</u> myocardial infarction	<u>VORSCHRIFTEN FÜR</u> den Myokardinfarkt
trigger in the translation is outside of the entity	She did <u>not</u> notice any blood / urine / emesis / stool in the bed.	Sie bemerkte <u>kein</u> Blut / Urin / Erbrechen / Stuhl im Bett.
missing of a possible trigger translation in NegEx-Ger	<u>Denies</u> fevers, pleuritic chest pain or cough.	Verleugnet Fieber, pleuritische Brustschmerzen oder Husten.
missing of translated compounds of type Entity + trigger in NegEx-Ger	She was pain free on the day of discharge.	Sie war am Tag der Entlassung schmerzfrei.
missing trigger phrase in NegEx-Ger due to word order	He then presented to Mass. Mental Health Center where he ruled out for an myocardial infarction by enzymes and electrocardiograms.	Er überreichte dann der Messe. Mental Health Center, wo er für einen Myokardinfarkt durch Enzyme und Elektrokardiogramme ausgeschlossen wurde.
different encoding of umlauts in text and NegEx-Ger	the hypernatremia fully resolved when he resumed eating on his own and had access to free water.	Die Hypernatrimie vollständig aufgelöst, als er wieder essen auf eigene Faust und hatte Zugang zu freien Wasser.

Table 7: Examples of the potential causes for prediction errors. The analysis focuses on the translation problems and the differences between the German and English language. The tags <E></E> enclose the entities, the factuality triggers are underlined. The original English examples originate from the i2b2 data.

- FactualMedBERT-DE on the translated i2b2 data and BRONCO150-BERT on the BRONCO150 data - this is no surprise. Moreover, the results indicate that the mixed model (i2b2+BRONCO150-BERT) performs generally well on all datasets, therefore might be the model of choice. However, it is important to note, that BRONCO150 has got an unusual label distribution. While *affirmed* is the most frequent label in all datasets, BRONCO has got an unusually high frequency of *possible* labels, which is connected to the way labels were mapped to the three final actuality labels. However, this might influence the actuality classification of other datasets.

## A.2 BERT Setup

For BERT, we used epochs number of 3/4 (for English and German BERT, respectively), a batch size

of 32, a dropout rate of 0.1, and a learning rate of  $1e - 5$ .

## A.3 Examples of Linguistic Error Analysis

Our analysis focuses on the potential sources for false predictions, in particular on causes related to the translation or the differences in the features of the German and English languages. Table 7 presents full-text examples from the original and translated data. For a detailed description of the possible issues see Section 4.1.