

Sebis at SemEval-2023 Task 7: A Joint System for Natural Language Inference and Evidence Retrieval from Clinical Trial Reports

Juraj Vladika and Florian Matthes

Department of Computer Science, Technical University of Munich
{juraj.vladika, matthes}@tum.de

Abstract

With the increasing number of clinical trial reports generated every day, it is becoming hard to keep up with novel discoveries that inform evidence-based healthcare recommendations. To help automate this process and assist medical experts, NLP solutions are being developed. This motivated the SemEval-2023 Task 7, where the goal was to develop an NLP system for two tasks: evidence retrieval and natural language inference from clinical trial data. In this paper, we describe our two developed systems. The first one is a pipeline system that models the two tasks separately, while the second one is a joint system that learns the two tasks simultaneously with a shared representation and a multi-task learning approach. The final system combines their outputs in an ensemble system. We formalize the models, present their characteristics and challenges, and provide an analysis of achieved results. Our system ranked 3rd out of 40 participants with a final submission.

1 Introduction

Clinical trials are research studies carried out in human subjects to assess the effectiveness and safety of medical, surgical, or behavioral interventions (Friedman et al., 2015). These investigations constitute the main approach that medical researchers use to determine whether a novel treatment such as a new medication, procedure, diet, or medical device is safe and efficient in humans. Clinical trials are often used to compare the effectiveness of a new treatment against the standard treatment or placebo treatment and to assess its adverse effects. When performed rigorously, clinical trials present the most valuable resource for informing evidence-based healthcare decisions.

On average, more than 100 reports of CTs are published every day (Zarin et al., 2019). Keeping up with all their results and novel discoveries is a time-consuming and often practically impossible endeavor. This has brought to light dedicated

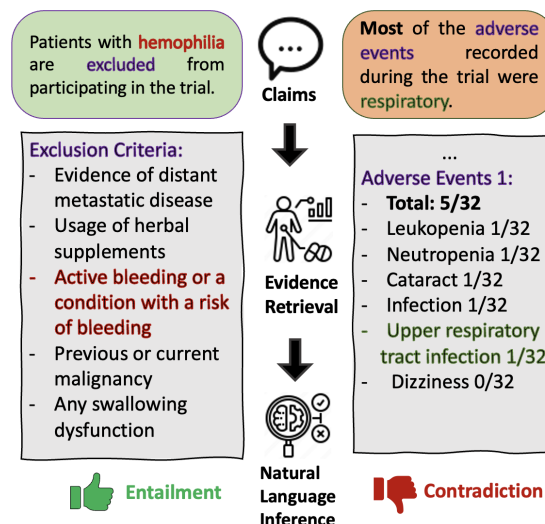


Figure 1: The task consists of predicting whether a given claim entails or contradicts the clinical trial report based on the evidence found in it.

organizations, such as Cochrane,¹ that manually synthesize these clinical outcomes (Higgins et al., 2019). Nevertheless, they struggle to keep up with the ever-increasing amount of literature and new studies. For this purpose, automated approaches based on Machine Learning (ML) and Natural Language Processing (NLP) can be developed to facilitate the process of inferring new knowledge and finding evidence from clinical trial reports.

The SemEval-2023 Task 7, titled *Multi-evidence Natural Language Inference for Clinical Trial Data (NLI4CT)*, focused on developing NLP systems for making conclusions and finding evidence in clinical trial reports (CTRs) (Jullien et al., 2023). It featured two subtasks that were strongly coupled together. The first subtask was, given a clinical trial document and a claim, to develop a model that infers a logical relation between them, namely either ENTAILMENT or CONTRADICTION. The second subtask was, given all the sentences in the clinical

¹<https://www.cochranelibrary.com>

trial document, to develop a model that selects a subset of those sentences that serve as evidence for making a decision on the logical entailment relation between the claim and the document. The task is illustrated in Figure 1. It shows two different claims, each coupled with an excerpt from a clinical trial report and a gold final label. The gold evidence sentences are highlighted for emphasis.

For this task, we developed two competitive NLP systems. The first one is a pipeline system, which learns to perform the tasks of evidence retrieval and textual entailment separately and takes into account one sentence at a time during the evidence selection process. The second approach is a joint system, which jointly learns the two subtasks with a unified representation of a claim and the whole document, and uses a shared training procedure. The joint system performed better on the test set, but the final system is an ensemble system that utilizes the power of both systems by combining their outputs with an averaging function. Our final system achieved the 3rd place in the competition out of 40 submitting participants. We provide the code of our system in a GitHub repository.²

2 Related Work

The task of Natural Language Inference (NLI) consists of inferring whether there is logical entailment or contradiction between two pieces of text – a premise and a hypothesis. It has been researched since the early days of NLP when it was mostly tackled with rule-based and linguistically informed approaches (MacCartney, 2009). A big surge of interest in the task occurred after the release of the SNLI dataset (Bowman et al., 2015), as well as follow-up datasets like MultiNLI (MNLI) (Williams et al., 2018) and Adversarial-NLI (ANLI) (Nie et al., 2020). The task has also been researched in the clinical domain – the dataset MedNLI (Romanov and Shivade, 2018) features more than 14,000 pairs of claims from clinical notes and patient records, extracted from the MIMIC-III database (Johnson et al., 2016).

The task of evidence retrieval, or more precisely evidence sentence selection, has been studied as a step in common NLP tasks like Machine Reading Comprehension (MRC) (Wang et al., 2019) and Question Answering (QA) (Thayaparan et al., 2021). Its main purpose is to improve the explainability and interpretability of models' predictions

but also to aid its reasoning process for making decisions. Combining the steps of evidence retrieval and textual entailment recognition comprises the task of automated fact-checking or claim verification, where the goal is to assess a claim's veracity based on the evidence related to it (Guo et al., 2022). This task has mostly been concerned with verifying claims related to news events, society, and politics, but has been increasingly so researched in scientific, biomedical, and clinical domains (Wadden et al., 2020; Sarrouiti et al., 2021).

3 Dataset

The clinical trials used for constructing the dataset originate from *ClinicalTrials.gov*,³ a database of privately and publicly funded clinical studies conducted around the world, run by the U.S. National Library of Medicine. All the clinical trial reports in the dataset are related to breast cancer and are written in English. There is a total of 1,200 clinical trials and 2,400 corresponding claims related to these CTRs, which were written and annotated by clinical domain experts, clinical trial organizers, and research oncologists.

For the task, each CTR may contain 1-2 patient groups, called cohorts or arms, and these groups may receive different treatments, or have different baseline characteristics. The dataset consists of a total of 2,400 claims and was split into a training dataset of 1,700 claims, a development dataset of 200 claims, and a hidden test dataset of 500 claims. There are two types of claims, those related only to a single CTR and other ones related to two CTRs, which usually make some form of comparison between the two reports. Each of the claims is related to only one of the following four sections:

- **Eligibility Criteria.** A set of conditions for patients to be allowed to take part in the clinical trial, such as age, gender, medical history.
- **Intervention:** Information concerning the type, dosage, frequency, and duration of treatments being studied in the clinical trial.
- **Results:** Reports the outcome of the trial with data like the number of participants, outcome measures, units, and results.
- **Adverse Events:** These are any unwanted side effects, signs, or symptoms observed in patients during the clinical trial.

²<https://github.com/jvladika/NLI4CT>

³<https://clinicaltrials.gov>

4 System Description

In this chapter, we will describe the architecture of the developed systems and the choice of base models comprising the systems. Even though the task features two standalone subtasks and participating in each of them was optional, we found the best synergy is achieved by combining them. The selection of appropriate evidence sentences in subtask 2 is an important prerequisite for recognizing textual entailment in subtask 1. This is because most textual entailment datasets used for training the NLI systems are modeled to recognize entailment between two sentences only, so the performance between a sentence and a whole document is still underwhelming (Schuster et al., 2022). Therefore, it was important to narrow down the whole document to a shorter span of text with only the relevant evidence sentences.

4.1 Pipeline Systems

In the pipeline system, the evidence sentences selected by the first model are used as the input for veracity prediction to the next model. Using the standard terminology from computer science, we call these systems pipeline systems. It is also common to use the same underlying base model in both steps and fine-tune it for these two different tasks (DeYoung et al., 2020).

To formalize the approach, we will define it mathematically. Given a claim c and n sentences s_1, s_2, \dots, s_n that constitute the clinical trial report, the goal is to train a model that predicts

$$z_i = \mathbb{1}[s_i \text{ is an evidence sentence}].$$

This is modeled as a binary sequence classification task, where candidate sequences are a concatenation of a candidate sentence s_i and the claim c in the form of $a_i = [s_i; SEP; c]$. Each of these sequences is encoded with the base language model to obtain their dense representation $h_i = BERT(a_i)$. This representation is then fed to a classifier model, Multi-Layer Perceptron (MLP), and its output is passed to a softmax function that assigns the probabilities on whether the candidate sentence is or is not the evidence:

$$p_i, \bar{p}_i = softmax(MLP(h_i)).$$

Finally, a selector function labels those sentences with a probability over a threshold to be evidence sentences – this threshold is a hyperparameter that

can be learned, such as $z_i = p_i > 0.5$. In the end, the model has selected k final evidence sentences e_1, e_2, \dots, e_k that are used as input for the next step.

The task of textual entailment is again a binary classification task that for a given claim c and k evidence sentences e_1, e_2, \dots, e_k predicts the logical relation of either ENTAILMENT or CONTRADICTION. Using the standard terminology of Natural Language Inference, the claim c is the hypothesis, and the concatenation of evidence sentences $e = [e_1; e_2; \dots; e_k]$ is the premise. These two are concatenated as $x = [c; SEP; e]$ and once again passed to the base language model to obtain the dense embedding $w = BERT(x)$. The model has to learn the function

$$\hat{y}(c; e) = softmax(MLP(w)),$$

which is the probability distribution of each inference label for the claim c given evidence e . The class with the highest probability score is selected as the final verdict $v(c; e) = argmax(y)$.

4.2 Joint Systems

The second system we developed jointly learns both the tasks of evidence retrieval and textual entailment. This leverages the machine-learning technique of multi-task learning (MTL), which was shown to be data efficient and improve the performance of each individual task it learns on through shared representations (Crawshaw, 2020). Intuitively, the model improves the performance on both tasks simultaneously since selecting high-quality evidence is important for recognizing entailment and conversely, the final entailment/contradiction label influences the specific evidence to be selected.

Unlike in the pipeline system, where each sequence consisted only of *one* candidate sentence and the claim, here the claim c is concatenated together with *all* of the sentences s_1, s_2, \dots, s_n in the clinical trial document to obtain a claim-document sequence $seq = [c; SEP; s_1; SEP; s_2; \dots; SEP; s_n]$.⁴ This approach makes the representation of each candidate sentence aware of the context it appears in with regard to the rest of the document, as well as aware of the claim itself. This sequence is fed to the base language model to obtain a dense representation $h =$

⁴Considering the input sequence length limitation of large language models, this sequence has to be truncated to 512 or 1024 tokens, which is good enough for the vast majority of documents in the dataset.

Base Model	Evidence Retrieval			Textual Entailment		
	Precision	Recall	F1	Precision	Recall	F1
BERT	77.5	80.8	80.2	58.2	64.0	61.0
BioBERT	78.0	81.9	80.8	64.3	65.0	64.5
ClinicalBERT	81.0	81.2	81.1	52.7	87.0	65.7
UmlsBERT	78.6	88.5	83.9	56.1	83.0	67.0
StructBERT	77.1	86.5	82.4	50.3	91.0	64.8
ERNIE	84.5	81.2	84.8	53.0	88.0	66.2
DeBERTa-v3	81.8	89.1	86.2	78.5	84.0	80.5

Table 1: Results of different base models for the two tasks on the development set.

$BERT(seq) = [h_c; SEP; h_{s_1}; \dots : SEP; h_{s_n}]$. The representation of each candidate sentence $h_{s_i} = [h_{w_1}, h_{w_2}, \dots, h_{w_m}]$ is singled out from the initial representation and passed to a binary linear classifier that, similarly to the one in the pipeline system, calculates the probabilities of the sentence being an evidence sentence:

$$p_i, \bar{p}_i = \text{softmax}(MLP(h_{s_i})).$$

Those sentences that are above the 0.5 threshold are then selected as evidence sentences and concatenated together to form the final evidence representation $h_e = [h_{e_1}, h_{e_2}, \dots, h_{e_k}]$. This representation is given to a final ternary linear classifier, that same as in the pipeline system predicts the verdict

$$v = \text{argmax}(\text{softmax}(MLP(h_e))).$$

Note that the representation of the claim h_c is not passed anywhere because the idea is for evidence sentences to already be aware of the semantics of the claim from the joint claim-document representation.

4.3 Ensemble System

After developing the pipeline system and the joint system, we ended up with the best-performing systems on the two tasks in each category of systems. These best systems were singled out and their outputs were combined for the final system. This type of approach is usually called a stacked (Pavlyshenko, 2018) or an ensemble system (Ganaie et al., 2021) and has been proven to perform well in machine-learning shared tasks and competitions.

Considering that the outputs of the two systems are different in certain predictions because of different output probabilities of classes, we decided to average the probabilities of each of the systems

with appropriate weights. After experimenting with different weights, the final function was:

$$p_{final} = 0.4 \cdot p_{pipeline} + 0.6 \cdot p_{joint}.$$

4.4 Base Models

Both of the components constituting the system use an underlying base model. We opt for large pre-trained language models (PLMs) since they represent the state of the art in virtually all NLP tasks. We experimented with a number of different base models. BERT (Devlin et al., 2019) is used as the representative vanilla PLM, which gives good initial insight into the performance of PLMs on the task. Over the years, there have been multiple domain-specific variations of the BERT model, specialized for text in the scientific, biomedical, or clinical domains. BioBERT (Lee et al., 2020) was fine-tuned on abstracts of biomedical scientific publications. ClinicalBERT (Alsentzer et al., 2019) is a model fine-tuned on MIMIC-III database, a database of electronic health records and clinical notes of patients admitted to critical care units. UmlsBERT (Michalopoulos et al., 2020) moves further away from unstructured text and injects structured domain knowledge, from UMLS (Bodenreider, 2004) – a large knowledge base of biomedical concepts and semantic relations between them, into the model training process.

DeBERTa (He et al., 2021b) is an improved extension of the BERT model that introduced disentangled attention and enhanced masked decoder, which both amplify the importance of positional embedding of tokens in a sequence. The DeBERTa-v3 (He et al., 2021a) is a novel version of the model that uses the task of replaced token detection (RTD) instead of predicting masked tokens. It was the first model to beat the human performance on the GLUE benchmark (Wang et al., 2018) for natural language understanding tasks. For compari-

son, we include other models that excel in NLU tasks, namely StructBERT (Wang et al., 2020) and ERNIE (Zhang et al., 2019). In the end, we would like to test whether the models highly specialized for inference and textual entailment tasks beat the BERT extension models fine-tuned to biomedical language and knowledge.

The hyperparameters were the same for all models and setups: learning rate 10^{-5} , warmup rate 0.06, weight decay 0.01, epochs 5–7, mixed precision training enabled. For all of the datasets, their *Large* version was used, imported from the HuggingFace repository.⁵ The models were additionally fine-tuned on the previously mentioned NLI datasets like MNLI, ANLI, MedNLI, following the approach of Laurer et al. (2022).

5 Results & Analysis

In this section, we report on results achieved by our systems and provide a qualitative error analysis with challenging examples from the dataset.

5.1 Final Results

The results of different base models for the two tasks with the pipeline system are presented in Table 1. These are the results on the dev set considering that this dataset has the revealed truth labels and allowed for an unlimited number of test runs. The results for the evidence retrieval are binary classification metrics, where each candidate sentence was of positive class if it constituted evidence, otherwise of negative class. The results of the textual entailment task are again binary classification metrics for the two classes of entailment and contradiction, using the gold evidence sentences. We found that the model trained on gold evidence performs better on the final test set than the model trained on internally selected evidence sentences.

As visible from Table 1, the NLI task of textual entailment recognition was more challenging than the evidence sentence selection task. The domain-specific biomedical and clinical BERT models outperformed the vanilla BERT model on both tasks, which shows the benefit of pre-training large language models on specialized text for NLP tasks concerning specific domains. The clinical BERT outperforms the BioBERT, which is an expected outcome since it works with clinical trial reports, while UmlsBERT outperforms both, showing the

⁵For example, DeBERTa-v3-Large is at: <https://huggingface.co/microsoft/deberta-v3-large>

System	Evidence F1	Entailment F1
Best pipeline	79.8	77.2
Best joint	80.4	78.3
Best ensemble	81.8	79.8
Place	4 th	3 rd

Table 2: Best performing model results on the test set

effectiveness of injecting structured knowledge into language models. Nevertheless, by far the best-performing base model turned out to be DeBERTa-v3 and it especially excelled in the textual entailment task. Although this was expected considering its proven efficiency on NLU tasks, the sheer margin of difference came as a surprise.

The joint systems turned out to perform slightly better on the evidence retrieval task. This was expected considering that it considers the full document when doing the evidence retrieval task, which means each candidate sentence is contextualized with appropriate surroundings around it and was already shown to perform well in evidence selection for fact-verification datasets (Stammach, 2021). Considering the long training times of the joint system, since it uses dense representations of full documents which can be up to 1024 tokens, we narrowed it down to just using DeBERTa as the underlying base model for embedding generation. The scores of the best-performing pipeline and best-performing joint system on the hidden test set are shown in Table 2. The outputs of these two systems were averaged with appropriate weights to obtain the final ensemble submission. We performed some additional post-processing such as truncating the selected evidence to a maximum of 20 sentences. This was done due to the tendency of the model to select all the candidate sentences as evidence sentences for claims with quantifiers like *at least one*, *exactly one*, *none of the patients*, etc. Selecting all the sentences incorrectly diminishes precision while increasing recall, but achieving high precision contributed more to the final score and was more challenging in general.

5.2 Error Analysis

To better understand the strengths and weaknesses of our system, we conducted a manual analysis

Challenge	Claim	Evidence	Label
Commonsense Reasoning	In order to participate in the trial, participants must be aware of where they are, and what day it is.	Inclusion Criteria: Cognitively oriented to time, place, and person (determined by nurse)	Entailment
Numerical Reasoning	Neutropenia affected the majority of patients in cohort 1 of the primary trial.	Adverse Events: Total 26/69 (37.68%) [...] Neutropenia 4/69 (5.80%)	Contradiction
Multi-Hop Reasoning	The primary trial and the secondary trial do not use the same route of administration for interventions.	(<i>primary</i>) Intervention: Vaccine Therapy / (<i>secondary</i>) Intervention: PET Guided Biopsy	Entailment
Medical Knowledge	Patients with cancer that has spread from a breast tumor to their CNS are able to take part in the trial.	No histologically proven bone marrow metastasis.	Contradiction
World Knowledge	A minimum bodyweight of 50kg is required to participate in the secondary trial.	Patient Characteristics: Total body weight > 110 lbs (without clothes)	Entailment

Table 3: Five common challenges found in the dataset with a representative example for each challenge that was classified incorrectly by our system.

of predictions on the development set to find out where the final model incorrectly labeled the claims as the opposite class. Table 3 provides five such examples from the dataset. Each example in the table consists of a claim, accompanying evidence sentences, and a gold label assigned to them by the expert annotators. We also identify five different challenges that the model has to tackle in order to be able to correctly classify these claims. These challenges are:

- **Commonsense Reasoning.** The model has to understand concepts and make judgments about everyday matters, that are common and innate to humans.
- **Numerical Reasoning.** The model has to be capable of applying basic mathematical operations and grasping comparisons, orders, or quantities.
- **Multi-hop Reasoning.** In some examples, the model needs to combine multiple pieces of evidence (make multiple "hops") to come to the final conclusion.
- **Medical Knowledge.** Certain claims rely on expert knowledge of medical concepts and terminology related to the human body, diseases, drugs, or treatments.
- **World Knowledge.** This type of knowledge refers to the non-linguistic information about the outside world contained in claims.

6 Conclusion

In this paper, we describe our solution for the SemEval-2023 Task 7, dealing with natural language inference and evidence retrieval from clinical trial reports. We motivate the task, discuss related work, provide formal definitions of the developed systems, present the results, analyze the performance of models, and discuss some challenges in the process. We developed two types of systems – a pipeline system, which learns evidence retrieval and textual entailment sequentially, and a joint system, which learns the two tasks simultaneously. The final system combines them into an ensemble and achieved the 3rd place in the competition out of 40 teams with a final submission.

We anticipate this system will be useful for facilitating the work of medical experts in synthesizing the results and outcomes of the ever-increasing amount of clinical trial reports, as well as be used by the NLP community for the related tasks of biomedical question answering, automated claim verification, and recognizing textual entailment. The system could be improved in the future by overcoming the challenges related to commonsense, numerical, and multi-hop reasoning, or by injecting additional medical and world knowledge to it.

References

Emily Alsentzer, John R Murphy, Willie Boag, Weihung Weng, Di Jin, Tristan Naumann, WA Redmond,

- and Matthew BA McDermott. 2019. Publicly available clinical bert embeddings. *NAACL HLT 2019*, page 72.
- Olivier Bodenreider. 2004. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl_1):D267–D270.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Michael Crawshaw. 2020. Multi-task learning with deep neural networks: A survey. *ArXiv*, abs/2009.09796.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020. [ERASER: A benchmark to evaluate rationalized NLP models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4443–4458, Online. Association for Computational Linguistics.
- Lawrence M Friedman, Curt D Furberg, David L DeMets, David M Reboussin, and Christopher B Granger. 2015. *Fundamentals of clinical trials*. Springer.
- Mudasir A Ganaie, Minghui Hu, et al. 2021. Ensemble deep learning: A review. *arXiv preprint arXiv:2104.02395*.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. [A survey on automated fact-checking](#). *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021a. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#). *CoRR*, abs/2111.09543.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021b. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Julian PT Higgins, James Thomas, Jacqueline Chandler, Miranda Cumpston, Tianjing Li, Matthew J Page, and Vivian A Welch. 2019. *Cochrane handbook for systematic reviews of interventions*. John Wiley & Sons.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9.
- Mael Jullien, Marco Valentino, Hannah Frost, Paul O’Regan, Donal Landers, and André Freitas. 2023. Semeval-2023 task 7: Multi-evidence natural language inference for clinical trial data. In *Proceedings of the 17th International Workshop on Semantic Evaluation*.
- Moritz Laurer, Wouter van Atteveldt, Andreu Casas, and Kasper Welbers. 2022. [Less annotating, more classifying – addressing the data scarcity issue of supervised machine learning with deep transfer learning and bert-nli](#). *Open Science Framework*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Bill MacCartney. 2009. *Natural language inference*. Stanford University.
- George Michalopoulos, Yuanxin Wang, Hussam Kaka, Helen Chen, and Alexander Wong. 2020. Umlsbert: Clinical domain knowledge augmentation of contextual embeddings using the unified medical language system metathesaurus. *arXiv preprint arXiv:2010.10391*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Bohdan Pavlyshenko. 2018. Using stacking approaches for machine learning models. In *2018 IEEE Second International Conference on Data Stream Mining & Processing (DSMP)*, pages 255–258. IEEE.
- Alexey Romanov and Chaitanya Shivade. 2018. [Lessons from natural language inference in the clinical domain](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596, Brussels, Belgium. Association for Computational Linguistics.
- Mourad Sarrouiti, Asma Ben Abacha, Yassine Mrabet, and Dina Demner-Fushman. 2021. [Evidence-based fact-checking of health-related claims](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3499–3512, Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Tal Schuster, Sihao Chen, Senaka Buthpitiya, Alex Fabrikant, and Donald Metzler. 2022. [Stretching sentence-pair NLI models to reason over long documents and clusters](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 394–412, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Dominik Stammach. 2021. [Evidence selection as a token-level prediction task](#). In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, pages 14–20, Dominican Republic. Association for Computational Linguistics.
- Mokanarangan Thayaparan, Marco Valentino, Peter Jansen, and Dmitry Ustalov. 2021. [TextGraphs 2021 shared task on multi-hop inference for explanation regeneration](#). In *Proceedings of the Fifteenth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-15)*, pages 156–165, Mexico City, Mexico. Association for Computational Linguistics.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. [Fact or fiction: Verifying scientific claims](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550, Online. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [Glue: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355.
- Hai Wang, Dian Yu, Kai Sun, Jianshu Chen, Dong Yu, David McAllester, and Dan Roth. 2019. [Evidence sentence extraction for machine reading comprehension](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 696–707, Hong Kong, China. Association for Computational Linguistics.
- Wei Wang, Bin Bi, Ming Yan, Chen Wu, Jiangnan Xia, Zuyi Bao, Liwei Peng, and Luo Si. 2020. [Structbert: Incorporating language structures into pre-training for deep language understanding](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). 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 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Deborah A Zarin, Kevin M Fain, Heather D Dobbins, Tony Tse, and Rebecca J Williams. 2019. 10-year update on study results submitted to clinicaltrials. gov. *New England Journal of Medicine*, 381(20):1966–1974.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. [ERNIE: Enhanced language representation with informative entities](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451, Florence, Italy. Association for Computational Linguistics.