

DKE-Research at SemEval-2024 Task 2: Incorporating Data Augmentation with Generative Models and Biomedical Knowledge to Enhance Inference Robustness

Yuqi Wang^{1,3}, Zeqiang Wang⁴, Wei Wang¹, Qi Chen¹,
Kaizhu Huang², Anh Nguyen³, Suparna De⁴

¹Xi'an Jiaotong Liverpool University ²Duke Kunshan University ³University of Liverpool ⁴University of Surrey
yuqi.wang17@student.xjtlu.edu.cn, {wei.wang03, qi.chen02}@xjtlu.edu.cn,
kaizhu.huang@dukekunshan.edu.cn, anh.nguyen@liverpool.ac.uk,
{zeqiang.wang, s.de}@surrey.ac.uk

Abstract

Safe and reliable natural language inference is critical for extracting insights from clinical trial reports but poses challenges due to biases in large pre-trained language models. This paper presents a novel data augmentation technique to improve model robustness for biomedical natural language inference in clinical trials. By generating synthetic examples through semantic perturbations and domain-specific vocabulary replacement and adding a new task for numerical and quantitative reasoning, we introduce greater diversity and reduce shortcut learning. Our approach, combined with multi-task learning and the DeBERTa architecture, achieved significant performance gains on the NLI4CT 2024 benchmark compared to the original language models. Ablation studies validate the contribution of each augmentation method in improving robustness. Our best-performing model ranked 12th in terms of faithfulness and 8th in terms of consistency, respectively, out of the 32 participants.

1 Introduction

In the domain of clinical trial analysis, researchers and practitioners are overwhelmed with an ever-expanding corpus of clinical trial reports (CTRs). The current repository contains a vast number of documents and is rapidly growing, a trend that correlates with the increasing prevalence of cross-national, cross-ethnic, and multi-center clinical studies (Bastian et al., 2010). This growth necessitates a scalable approach to evaluate and interpret the massive amount of data in these reports (Goldberg et al., 2017; Li and Bergan, 2020).

Recent advances in Natural Language Processing (NLP) offer promising avenues for the automated analysis of CTRs. Such analyses include medical evidence understanding (Nye et al., 2021), information retrieval (Wang et al., 2023b), causal relationship identification (Cai et al., 2017), and

the inference of underlying reasons for trial outcomes (Steinberg et al., 2023). Integrating natural language inference (NLI) with CTRs has the potential to revolutionize the large-scale, NLP-based examination of experimental medicine (Kim and Delen, 2018). Despite the progress in NLP, the application of large language models to this task presents several challenges, including susceptibility to shortcut learning, hallucination, and biases stemming from word distribution patterns within the training data (Huang et al., 2023).

To address these issues, we propose a novel method that leverages generative language models, such as GPT-3.5¹, and biomedical domain knowledge graphs to enhance data diversity. Our approach introduces three types of data augmentation: numeric question-answering data generation, semantic perturbations, and domain-tailored lexical substitutions for the biomedical field. By combining these data augmentation techniques with multi-task learning and the DeBERTa (He et al., 2021) architecture, we have achieved significant improvements in terms of faithfulness and consistency on the NLI4CT 2024 dataset. This paper outlines our approach, elaborates on the design of the perturbations and the multi-task learning process, and demonstrates the efficacy of our method through rigorous evaluation.

2 Background

In a crucial field like healthcare, where misinterpretations can have severe implications, NLI models must present precise predictions and reliable interpretations. This highlights the importance of accurate and trustworthy reasoning in these NLI models.

SemEval 2024 Task 2 (Jullien et al., 2024) provides multi-sentence textual data consisting of patient case histories and medical reports. The objec-

¹<https://openai.com/chatgpt>

tive of this task is to predict the logical relationship between the CTR and a given statement, including entailment and contradiction. The evaluation emphasizes prediction accuracy as well as the robustness to the controlled interventions, helping increase healthcare practitioners’ trust in the system’s predictions.

Enhancing the robustness of NLI models for healthcare can be strategically achieved using data augmentation techniques. Synthetic data generation via techniques like conditional text generation can expand training data diversity and volume to improve model generalization capabilities (Liu et al., 2020; Puri et al., 2020; Bayer et al., 2023). Meanwhile, multi-task learning with auxiliary objectives related to logical reasoning and explanation generation can enhance faithful reasoning abilities (Li et al., 2022). Useful domain knowledge can be captured by training language models on domain-specific medical textual datasets (Singhal et al., 2023; Tian et al., 2024). Complementary data-centric methods can augment model architecture design to develop more capable, trustworthy, and clinical NLI systems.

3 System overview

In this section, we describe the proposed system to tackle the NLI problem and enhance the model’s robustness against interventions spanning numerical, vocabulary, and semantic dimensions, as shown in Figure 1.

3.1 Data for Numerical Question Answering Task

A major limitation of many language models lies in their tendency to learn linguistic patterns and features from large-scale textual data while lacking capabilities for numerical and quantitative reasoning (Geva et al., 2020). Such capabilities are crucial for analyzing relationships between CTRs and corresponding claims. Although BERT-based models pre-trained on NLI tasks, i.e. DeBERTa, can conduct general linguistic inference, they remain vulnerable to numerical perturbations in statements.

Therefore, we propose to leverage GPT-3.5 to generate data tailored to the numerical question-answering task based on original entailed statements: The entailed statement, denoted as x , corresponding to a given CTR, is converted into a question q that requires numerical reasoning. Subsequently, three candidate choices c are enumer-

ated, each accompanied by an answer a extracted from the original statements. The loss function employed for this task is binary cross-entropy and is expressed as follows:

$$\mathcal{L}_{NQA} = \begin{cases} -\log g_{\theta}(\text{CTR}, q, c;) & c = a \\ -[1 - \log g_{\theta}(\text{CTR}, q, c)] & c \neq a \end{cases}$$

where $g(\cdot)$ is the function to determine if the candidate choice is the correct answer, and θ is the corresponding parameters for the DeBERTa backbone network and the additional classifier.

This numerical question-answering task serves as an auxiliary task to enhance numerical reasoning abilities. The final loss function for the system combines the losses from this task and the main NLI task, i.e.

$$\mathcal{L} = \mathcal{L}_{NLI} + \lambda \mathcal{L}_{NQA}$$

where λ is the hyper-parameter to be tuned in the validation phase.

3.2 Semantic Perturbation

We utilize GPT-3.5 to generate perturbed statements based on the original entailed input, obtaining both semantic-altering variants labeled as “contradictions” and semantic-preserving variants labeled as “entailment”. Specifically, to produce contradictory versions, guiding keywords such as “contradicted” and “minor changes” are injected into the input prompt to slightly modify the original statement while altering the semantics to create a contradiction. Conversely, to generate entailed versions, guiding phrases such as “paraphrase” are included in the prompt to rephrase the statement extensively while retaining semantic equivalence. This controlled semantic perturbation of the input statement via guided text generation allows us to efficiently augment the dataset with both contradicting and entailing variants of the original input.

3.3 Vocabulary Replacement

When we analyze textual data in the clinical domain, we need to pay attention to the vocabulary because it contains many terms that are specific to this domain (Wang et al., 2018). However, most NLI models are pre-trained on data from general domains, and they are unaware of the meaning or relevance of these terms (Wang et al., 2023a). To address this problem, we use a combination of biomedical knowledge graph embedding and statistical model, which can help us find the most important keyword to replace the term in the statement

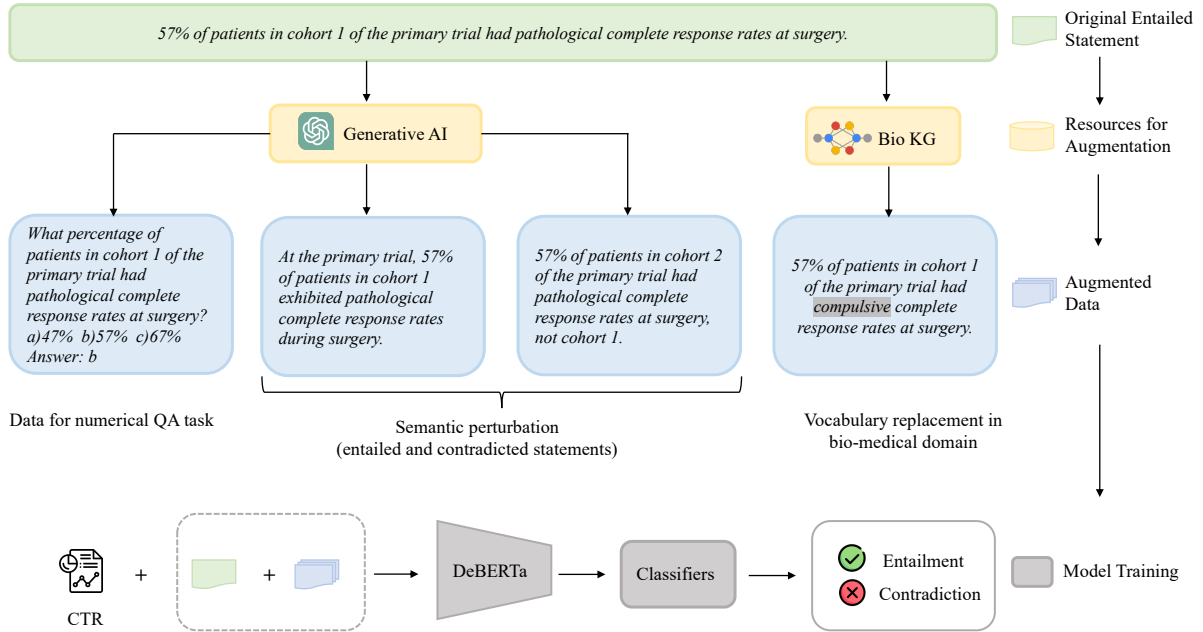


Figure 1: The overall demonstration of the proposed system. The upper part of the demonstration involves the application of data augmentation techniques to entailed statements extracted from the original NLI dataset, leveraging generative artificial intelligence (AI) and biomedical domain knowledge graphs. Specifically, we undertake the following procedures: 1) Transformation of statements into multiple-choice questions accompanied by corresponding answers; 2) Introduction of semantic perturbations to the original entailed statements; 3) Employing a statistical method to identify keywords within the original entailed statements, followed by their substitution with synonyms sourced from the biomedical knowledge graph. In the lower part of the demonstration, we incorporate the original entailed statements, augmented data, and CTRs as training data to develop a classifier based on the DeBERTa architecture.

and generate the augmented data to improve the vocabulary alignment. Specifically, given a statement x , consisting of n words, i.e. $x = \{w_1, w_2, \dots, w_n\}$ and the set of all the statements, denoted as D , we first remove all the stop-words and apply Term-Frequency-Inverse Document Frequency (TF-IDF) to identify the most important term in the statement, i.e.

$$w^* = \arg \max_{w_i \in x} \text{TF}(w_i, x) \times \text{IDF}(w_i, D)$$

Subsequently, we locate a term in the biomedical embedding space that shares the same part-of-speech and has the highest similarity score with the chosen term, using it as the substitute, i.e.

$$\hat{w}^* = \arg \max_{w \in V} \{ \text{sim}(w^*, w) \mid \text{PoS}(w) = \text{PoS}(w^*) \}$$

where V is the biomedical term vocabulary and $\text{PoS}(\cdot)$ is the part-of-speech of a word. In this way, we can substitute w^* in the original statement with \hat{w}^* to generate a new adversarial sample to enhance the model robustness in the vocabulary aspect.

4 Experimental setup

4.1 Dataset

	Ent.	Con.	Alt.	Pres.	SUM
Train	850	850	-	-	1,700
Val.	100	100	1,606	336	2,142
Test	250	250	4,136	864	5,500

Table 1: Statistics of the validation and test set. ‘‘Ent.’’ and ‘‘Con.’’ stands for entailment and contradiction, while ‘‘Alt.’’ and ‘‘Pres.’’ stands for altering and preserving.

We conducted experiments on the NLI4CT 2024 dataset (Jullien et al., 2024), generated by clinical domain experts and sourced from a large database for clinical studies². The statistic of this dataset is summarized in Table 1. The training data is the same as the NLI4CT 2023 dataset (Jullien et al., 2023) while there are perturbed samples in the validation and testing sets.

²<https://ClinicalTrials.gov>

4.2 Metrics

We first assessed the performance of the original statements without any perturbation and recorded the corresponding F1 score, precision, and recall. Then, we assessed the performance of the contrast set, consisting of interventions. Specifically, to evaluate the model’s robustness to the semantic-preserving interventions, we used consistency as the metric, i.e.

$$\text{Consistency} = \frac{1}{N} \sum_1^N 1 - |f(x'_i) - f(x_i)|$$
$$x'_i \in C : \text{Label}(x_i) = \text{Label}(x'_i)$$

Where C is the contrast set, and N is the number of the statements in the contrast set. x'_i is the perturbed statement for x_i and $f(\cdot)$ computes the final prediction from the model. For the semantic-altering interventions, we evaluated the model using faithfulness, i.e.

$$\text{Faithfulness} = \frac{1}{N} \sum_1^N |f(x'_i) - f(x_i)|$$

$$x'_i \in C : \text{Label}(x_i) \neq \text{Label}(x'_i), \text{ and } f(x_i) = \text{Label}(x_i)$$

4.3 Implementation details

We downloaded DeBERTa models from the Huggingface repository³ and implemented our proposed method based on Python 3.10 and Pytorch 2.1.1. During the model training, we used the Adam optimizer and set the learning rate to $5e - 6$ with a batch size of 4, following the original work (He et al., 2021). The maximum sequence length the model can take was set to 512. The epoch number was set to 20, and the early stopping based on the validation set was applied to avoid overfitting. The input format for the NLI task in this work is structured as follows: [CLS] + CTR + [SEP] + claim + [SEP]. In this structure, [CLS] serves as the initial token for classification in DeBERTa, and [SEP] acts as a separator token. For the vocabulary replacement, we used the bio-medical domain embedding from the work by (Zhang et al., 2019), which has been pre-trained over the MeSH knowledge graph⁴. For preprocessing, such as stop word filtering and part-of-speech tagging, we used the NLTK library⁵ in Python. We include prompts for numerical question-answering data generation and semantic perturbation in Table 2.

³<https://huggingface.co/>

⁴<https://www.ncbi.nlm.nih.gov/mesh/>

⁵<https://www.nltk.org/>

5 Results

We conducted experiments with different-sized DeBERTa models, iteratively adding augmented data from three different interventions to the training set. As shown in Table 3, incorporating all three types of augmented data greatly improved the average faithfulness and consistency scores. Specifically, we witnessed gains of 8.17% on DeBERTa-l and 2.37% on DeBERTa-b. This result also suggests that the augmented training data provided more benefit to the larger-sized DeBERTa model in terms of robustness. The additional augmented examples may have provided useful regularization, helping it generalize better on both the unaltered control and contrast datasets. Our best-performing model ranked 12th in terms of faithfulness and 8th in terms of consistency, respectively, out of the 32 participants.

From this iterative process, we can see that semantic perturbation with generative AI contributes mainly to the performance gain for both NLI models. Compared with this, vocabulary replacement in the biomedical domain has only a minor effect. This may suggest that vocabulary replacement in our work may be relatively less effective in this case because it only swaps out individual words, while semantic perturbation modifies the whole statement. Hence, semantic perturbation provides more meaningful variations to augment the training data.

While the augmented data improved the robustness to interventions, we noticed a slight performance drop in the control set. For example, the F1 score on the control set decreased by 3.16% for DeBERTa-l and 0.48% for DeBERTa-b after adding all the augmented data. This performance decline indicates there may have been a small trade-off between improving robustness to interventions and maintaining strong performance on the original data. One of the reasons accounting for this could be that the generative AI may generate noisy or irrelevant data. For example, in numerical question answering data generation, if the original entailed statement discusses an assumption about a 50-year-old patient not mentioned in the CTR, the generative model may create an unrelated question about the patient’s age that cannot be inferred from the given information. Another example involves vocabulary replacement: we observed that there exist some cases where even two words having very similar embeddings in the biomedical domain

	Prompt
NQA	<i>Please convert the statement to a multiple choice question that requires the numerical or quantitative reasoning, and each question has 3 choices, using the given template: \n Question: [Question] \n Choices: 1. [Choice 1] \n 2. [Choice 2] \n 3. [Choice 3] \n Correct Answer: [Correct Answer].</i>
SP.-Ent.	<i>Please rephrase the given statement:</i>
SP.-Con.	<i>Please generate a contradictory statement based on the given statement, with a minor change:</i>

Table 2: Prompts for numerical question-answering data generation and semantic perturbation. NQA stands for numerical question answering. SP.-Con. and SP.-Ent. means semantic perturbation to generate statements labeled as contradiction and entailment, respectively.

Method	Validation					Test				
	F1	Prec.	Rec.	Faith.	Con.	F1	Prec.	Rec.	Faith.	Con.
DeBERTa-l	81.82	90.00	75.00	73.81	71.48	77.25	80.80	73.99	67.13	71.06
+SP	81.77	83.00	80.58	85.42	75.16	75.52	72.80	78.45	78.24	74.01
+VR	81.00	81.00	81.00	86.01	74.16	75.05	71.60	78.85	78.59	74.42
+NQA	80.60	81.00	80.20	86.61	74.91	74.09	69.20	79.72	79.98	74.54
DeBERTa-b	70.87	73.00	68.87	49.40	60.02	62.53	60.40	64.81	57.75	59.33
+SP	71.84	74.00	69.81	51.49	60.65	62.08	59.60	64.78	60.65	59.70
+VR	70.59	72.00	69.23	52.38	60.71	62.21	59.60	65.07	60.76	59.72
+NQA	70.30	71.00	69.61	52.98	60.77	62.05	59.20	65.20	61.92	59.89

Table 3: Results on the development set and testing set for NLI4CT 2024 dataset. DeBERTa-l and DeBERTa-b are the large version and base version of the DeBERTa model, respectively. SP and VR stand for semantic perturbation and vocabulary replacement. The best results for F1 score on the control set, faithfulness, and consistency are highlighted.

knowledge graph embedding space may not be very closely related in the context of the current statement. Including these illogical examples in the augmented training data could mislead the original DeBERTa model, resulting in worse performance on the unaltered control set.

6 Conclusion

In this work, we proposed a data augmentation approach to enhance the robustness of natural language inference models for clinical trial report analysis. Our method leverages generative AI and biomedical knowledge graphs to augment training data along three dimensions: numerical reasoning, semantic perturbations, and domain-tailored lexical substitutions. Experiments on the NLI4CT 2024 dataset demonstrate that our approach effectively improves model faithfulness and consistency against controlled interventions, with significant

gains against the DeBERTa baselines.

However, we observed a slight performance drop on the unaltered test set, indicating a trade-off between robustness to perturbations and maintaining strong performance on original data. Future work will focus on: 1) generating higher-quality augmented examples using numerical question-answering data generation to minimize or avoid performance drop; 2) validating the perturbed samples to help remove noisy or irrelevant examples (Wang et al., 2023c); 3) incorporating external structured knowledge via pre-training on knowledge graphs and not just lexical substitution, which can provide more contextual domain information.

7 Acknowledgments

We would like to thank all the anonymous reviewers for their valuable feedback. We would like to acknowledge the financial support provided by the

Postgraduate Research Scholarship (PGRS) (contract number PGRS-20-06-013) at Xi'an Jiaotong-Liverpool University. Additionally, this research has received partial funding from the Jiangsu Science and Technology Programme (contract number BK20221260) and the Research Development Fund (contract number RDF-22-01-132) at Xi'an Jiaotong-Liverpool University.

References

- Hilda Bastian, Paul Glasziou, and Iain Chalmers. 2010. Seventy-five trials and eleven systematic reviews a day: how will we ever keep up? *PLoS medicine*, 7(9):e1000326.
- Markus Bayer, Marc-André Kaufhold, Björn Buchhold, Marcel Keller, Jörg Dallmeyer, and Christian Reuter. 2023. Data augmentation in natural language processing: a novel text generation approach for long and short text classifiers. *International journal of machine learning and cybernetics*, 14(1):135–150.
- Ruichu Cai, Mei Liu, Yong Hu, Brittany L Melton, Michael E Matheny, Hua Xu, Lian Duan, and Lemuel R Waitman. 2017. Identification of adverse drug-drug interactions through causal association rule discovery from spontaneous adverse event reports. *Artificial intelligence in medicine*, 76:7–15.
- Mor Geva, Ankit Gupta, and Jonathan Berant. 2020. Injecting numerical reasoning skills into language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 946–958.
- Richard M Goldberg, Lai Wei, and Soledad Fernandez. 2017. The evolution of clinical trials in oncology: defining who benefits from new drugs using innovative study designs.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*.
- Maël Jullien, Marco Valentino, and André Freitas. 2024. SemEval-2024 task 2: Safe biomedical natural language inference for clinical trials. In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*. Association for Computational Linguistics.
- Mael Jullien, Marco Valentino, Hannah Frost, Paul O'Regan, Dónal Landers, and Andre Freitas. 2023. [NLI4CT: Multi-evidence natural language inference for clinical trial reports](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16745–16764, Singapore. Association for Computational Linguistics.
- Yong-Mi Kim and Dursun Delen. 2018. Medical informatics research trend analysis: A text mining approach. *Health informatics journal*, 24(4):432–452.
- Allen Li and Raymond C Bergan. 2020. Clinical trial design: Past, present, and future in the context of big data and precision medicine. *Cancer*, 126(22):4838–4846.
- Wei Li, Wenhao Wu, Moye Chen, Jiachen Liu, Xinyan Xiao, and Hua Wu. 2022. Faithfulness in natural language generation: A systematic survey of analysis, evaluation and optimization methods. *arXiv preprint arXiv:2203.05227*.
- Ruibo Liu, Guangxuan Xu, Chenyan Jia, Weicheng Ma, Lili Wang, and Soroush Vosoughi. 2020. Data boost: Text data augmentation through reinforcement learning guided conditional generation. *arXiv preprint arXiv:2012.02952*.
- Benjamin E Nye, Jay DeYoung, Eric Lehman, Ani Nenkova, Iain J Marshall, and Byron C Wallace. 2021. Understanding clinical trial reports: Extracting medical entities and their relations. *AMIA Summits on Translational Science Proceedings*, 2021:485.
- Raul Puri, Ryan Spring, Mostofa Patwary, Mohammad Shoenybi, and Bryan Catanzaro. 2020. Training question answering models from synthetic data. *arXiv preprint arXiv:2002.09599*.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2023. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180.
- Ethan Steinberg, Nikolaos Ignatiadis, Steve Yadlowsky, Yizhe Xu, and Nigam Shah. 2023. Using public clinical trial reports to probe non-experimental causal inference methods. *BMC Medical Research Methodology*, 23(1):204.
- Shubo Tian, Qiao Jin, Lana Yeganova, Po-Ting Lai, Qingqing Zhu, Xiuying Chen, Yifan Yang, Qingyu Chen, Won Kim, Donald C Comeau, et al. 2024. Opportunities and challenges for chatgpt and large language models in biomedicine and health. *Briefings in Bioinformatics*, 25(1):bbad493.
- Yanshan Wang, Sijia Liu, Naveed Afzal, Majid Rastegar-Mojarad, Liwei Wang, Feichen Shen, Paul Kingsbury, and Hongfang Liu. 2018. A comparison of word embeddings for the biomedical natural language processing. *Journal of biomedical informatics*, 87:12–20.
- Yuqi Wang, Wei Wang, Qi Chen, Kaizhu Huang, Anh Nguyen, Suparna De, and Amir Hussain. 2023a. Fusing external knowledge resources for natural language understanding techniques: A survey. *Information Fusion*, 92:190–204.

Yuqi Wang, Zeqiang Wang, Wei Wang, Qi Chen, Kaizhu Huang, Anh Nguyen, and Suparna De. 2023b. Zero-shot medical information retrieval via knowledge graph embedding. In *International Workshop on Internet of Things of Big Data for Healthcare*, pages 29–40. Springer.

Zimu Wang, Wei Wang, Qi Chen, Qiufeng Wang, and Anh Nguyen. 2023c. Generating valid and natural adversarial examples with large language models. *arXiv preprint arXiv:2311.11861*.

Yijia Zhang, Qingyu Chen, Zhihao Yang, Hongfei Lin, and Zhiyong Lu. 2019. Biowordvec, improving biomedical word embeddings with subword information and mesh. *Scientific data*, 6(1):52.