Zero-shot sampling of adversarial entities in biomedical question answering

Anonymous ACL submission

Abstract

The increasing depth of parametric domain knowledge in large language models (LLMs) is fueling their rapid deployment in real-world applications. In high-stakes and knowledge-intensive tasks, understanding model vulnerabilities is essential for quantifying the trustworthiness of model predictions and regulating their use. The recent discovery of named entities as adversarial examples in natural language processing tasks raises questions about their potential guises in other settings. Here, we propose a distance-weighted sampling scheme in embedding space to discover diverse adversarial entities as distractors. We demonstrate its advantage over random sampling in adversarial question answering on biomedical topics. Our approach enables the exploration of different regions on the attack surface, which reveals two regimes of adversarial entities that markedly differ in their characteristics. Moreover, we show that the attacks successfully manipulate token-wise Shapley value explanations, which become deceptive in the adversarial setting. Our investigations illustrate the brittleness of domain knowledge in LLMs and reveal a shortcoming of standard evaluations for high-capacity models[1].

1 Introduction

LLMs such as pre-trained and finetuned generalist language models (GLMs) and their domain-adapted versions (Zhao et al., 2023; Wang et al., 2023a) incorporating transformer architectures are emerging as the technological backbone of next-generation search engines and knowledge bases (Petroni et al., 2019; Razniewski et al., 2021; Sung et al., 2021). They are increasingly used in knowledge-intensive tasks by the general public with little understanding of their inherent issues in trustworthiness and safety (Nat, 2023; Spitale et al., 2023; Si et al., 2023). Adversarial examples are valuable probes for model vulnerability and robustness, and are essential for performance evaluation and regulatory audits (Shayegani et al., 2023). For language models, adversarial examples are generated through guided search (Alzantot et al., 2018), sampling and probabilistic methods (Ren et al., 2019; Yang et al., 2020; Yan et al., 2022), or collected dynamically by human operators (Jia and Liang, 2017; Wallace et al., 2019b). While the search methods iteratively find random string adversarial examples (Wallace et al., 2019a), the sampling approach obviates the need for predetermined search heuristics. It can operate in the zero-shot setting and is therefore fitting for model auditing frameworks (Metaxa et al., 2021).

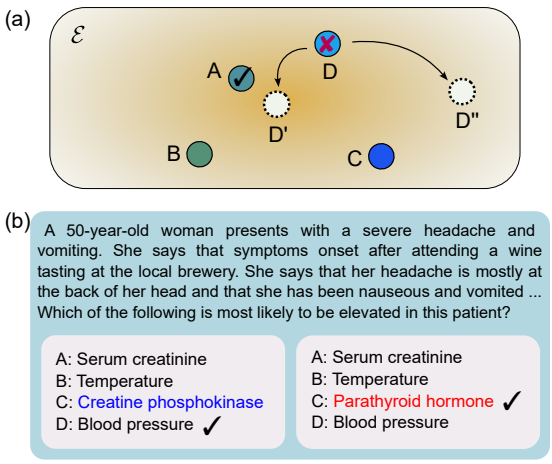

Figure 1: Text substitution attack on QA with adversarial distractors. (a) Illustrated attack scheme in embedding space. A is the key to the question, B-D the original distractors, D′ and D″ the adversarial distractors. (b) An example of the attack showing the correct (left) and incorrect (right) model responses (checkmarked) before and after perturbation.

The present work focuses on question answering (QA) in the biomedical domain, which has been the proving ground in recent achievements of biomedical language models (BLMs), here referring to GLMs adapted to the biomedical domain (Sing-

[1]The code and datasets developed for the work are available in an anonymized repository on GitHub.

hal et al., 2023a; Kung et al., 2023; Singhal et al., 2023b). Oftentimes, GLMs have already acquired a certain amount of general biomedical knowledge during model development. BLMs gain further domain knowledge primarily through finetuning on a variety of text data from clinical notes (Lehman et al., 2023), electronic health records (Wornow et al., 2023), to abstracts or excerpts from scholarly publications (Tinn et al., 2023). Although QA is frequently employed for model performance evaluation (Clark and Etzioni, 2016; Gardner et al., 2019; Robinson and Wingate, 2023), currently, the performance of LLMs in biomedical QA lacks systematic evaluation and their deployments in the real world are still in their infancy (Kung et al., 2023). Specifically, model performance against adversarial examples is seldomly assessed and their efficient sampling in the biomedical domain remains challenging, leaving a gap in our understanding of model robustness.

Knowledge and reasoning tests in the biomedical domain assess the understanding of diverse concepts, which are represented by domain-specific named entities (NEs). In this work, we develop an in-domain adversarial probe for model knowledge on biomedical NEs through text substitution of entity mentions in the distractors of multiple-choice QA. We introduce distance-weighted sampling to generate adversarial examples of distractors, or *adversarial distractors*, in the zero-shot setting. The attack scheme targets two of the most common NEs in the biomedical domain (Wei et al., 2019): drug and disease names. In addition, we introduce the `FDA-drugs` and `CTD-diseases` datasets as the source vocabularies of the perturbation sets, along with an entity type-annotated version of existing biomedical QA datasets MedQA-USMLE (Jin et al., 2021) and MedMCQA (Pal et al., 2022). Through analysis of attack statistics, we show that our approach provides more flexibility to uniform random sampling in entity-centric attacks (Yan et al., 2022). Varying the exponent of the distance-dependent probability weights allows tuning of the characteristics of adversarial entities, which maps the entity-level robustness landscapes of LLMs with different architectures, encoding types, and development history. We further explore the relationship between Shapley value-based model explanation (Chen et al., 2023) and adversarial attacks and uncover the defining signatures of successful attacks in the current context.

2 Entity-centric adversarial attack

2.1 Adversarial distractor generation

We focus on constructing adversarial distractors in multiple-choice QA. Let each question $\mathcal{Q}$ be represented by a tuple, $\mathcal{Q} = (\mathcal{C}, \mathcal{S}, \mathcal{D}, k)$, with the context $\mathcal{C}$, question stem $\mathcal{S}$, a distractor set $\mathcal{D} = \{d_i\}_{i=1}^{l}$, and a key k. The distractors are the wrong choices given to confuse test-takers, while the key is the correct one (CH and Saha, 2020). QA is a natural language generation task. In an autoregressive LLM, the answer corresponds to the next tokens following $(\mathcal{C}, \mathcal{S}, \mathcal{D}, k)$. The adversarial distractor is a threat model in which the perturbed distractor set $\mathcal{D}'$ takes the form

$$\mathcal{D}' = (\mathcal{D} \setminus \{d\}) \cup \{d'\}, \tag{1}$$

where d and d' are the original and perturbed distractors, respectively. The substitution is subject to $\|\mathcal{Q}' - \mathcal{Q}\| < \epsilon$, where ϵ is the perturbation budget and the norm $\|\cdot\|$ refers to edit distance. For simplicity, we assume the distractors only contain NEs and each distractor involves only one NE (such as in Fig. 1b). An adversarially perturbed question elicits a false response from the model, while it answers correctly to the unperturbed question,

$$\begin{aligned} &\operatorname{argmax} \Pr(Y = k | \mathcal{C}, \mathcal{S}, \mathcal{D}', k) \\ &\qquad \neq \operatorname{argmax} \Pr(Y = k | \mathcal{C}, \mathcal{S}, \mathcal{D}, k). \end{aligned} \tag{2}$$

The formulation in Eq. (2) is an untargeted attack, while the targeted version only requires restricting the answer to the perturbed question to a specific distractor. The perturbations are drawn from a finite perturbation set, denoted as $\mathcal{E}_{\text{-}k} := \mathcal{E} \setminus \{k\}$, with $\mathcal{E} = \{e_j\}_{j=1}^{m+1}$ being the entity dataset as illustrated in Fig. 1a. In practice, $\mathcal{E}$ may be compiled from potential model use cases or curated by regulatory bodies in a model audit.

We consider here distance-weighted sampling from the embedding space of NEs. A simple way to assign the weights is to use a power function. Specifically, for a key k, we assign a distance-dependent probability weight p_j to entity e_j in $\mathcal{E}_{\text{-}k}$ and the distractor d is substituted by

$$d' = e_j \quad \text{w.p.} \quad p_j = \frac{h^n(k, e_j)}{\sum_j h^n(k, e_j)}. \tag{3}$$

Here, $h(k, e_j)$ is a distance metric between k and e_j, while the exponent $n \in \mathbb{R}$ controls the shape of the probability weight distribution. When $n > 0$,

Table 1: Adversarial testing of model performance on subsets of MedQA-USMLE dataset[3]

	Drug-mention questions			Disease-mention questions		
Model (size)	NoPtb	RdPtb	DistPtb (-/+)	NoPtb	RdPtb	DistPtb (-/+)
Flan-T5-base (0.25B)	0.276	0.238	0.230 / 0.235	0.260	0.226	0.203 / 0.221
Flan-T5-large (0.78B)	0.287	0.252	0.245 / 0.243	0.294	0.249	0.230 / 0.224
Flan-T5-xl (3B)	0.318	0.285	0.275 / 0.284	0.335	0.298	0.281 / 0.293
Flan-T5-xxl (11B)	0.323	0.292	0.279 / 0.287	0.345	0.303	0.276 / 0.301
Flan-UL2 (20B)	0.382	0.349	0.335 / 0.342	0.396	0.356	0.334 / 0.344
GPT-3.5 (175B)	0.438	0.398	0.389 / 0.391	0.453	0.414	0.401 / 0.413
MedAlpaca-7B	0.460	0.431	0.421 / 0.416	0.488	0.439	0.423 / 0.436
Palmyra-Med-20B	0.463	0.413	0.405 / 0.414	0.500	0.441	0.425 / 0.441
PaLM2 (350B)	0.287	0.263	0.254 / 0.260	0.328	0.302	0.283 / 0.297

as n increases, the entities further away from k are more preferably weighted. When $n < 0$, the sampling method employs inverse distance weights. The deterministic limit[2] corresponds to $n \to +\infty$, when p_j is vanishingly small for all but the entity furthest from k in the embedding space according to $h(k, e_j)$,

$$d' = e_j \quad \text{s.t.} \quad e_j = \underset{e_j \in \mathcal{E}_{\text{-}k}}{\operatorname{argmax}}\, h(k, e_j). \quad (4)$$

When $n = 0$, p_j is the same for each entity in $\mathcal{E}_{\text{-}k}$, corresponding to uniform random sampling. A short proof of the limiting cases is given in Appendix A. Therefore, our approach bridges the random and deterministic approach to sampling adversarial examples. When $n = 0$, p_j is the same for each entity in $\mathcal{E}_{\text{-}k}$, corresponding to uniform random sampling. A short proof of the limiting cases is given in Supplement A. Therefore, our approach bridges the random and deterministic approach to sampling adversarial examples.

2.2 Datasets

Entity datasets To relate to real-world scenarios, we sourced the vocabulary datasets of drug and disease names from existing public databases. The drug names dataset (`FDA-drugs`) contains over 2.3k entities from known drugs approved by the United States Food and Drug Administration and curated by Drug Central[4] (Ursu et al., 2017). The disease names dataset[5] (`CTD-diseases`) contains over 9.8k entities from the Comparative Toxicogenomics Database (Davis et al., 2009). Minor data processing was carried out before the numerical experiments.

Biomedical QA datasets We selected over 9.3k questions from the MedQA-USMLE (Jin et al., 2021) dataset and over 3.8k questions from the MedMCQA (Pal et al., 2022) dataset for benchmarking. Both datasets are entity-rich and cover a wide range of biomedical topics. The MedQA-USMLE dataset contains long-context questions that are used to assess medical students in the United States. The MedMCQA dataset contains short questions used in medical exams in India. Both datasets are publicly available and don't contain personal information. We annotated the NEs in the two datasets according to entity types specified in the Unified Medical Language System (UMLS) (Bodenreider, 2004) using `scispaCy` (Neumann et al., 2019). We then generated the binary mention labels for the selected questions according to the UMLS entity types (see Appendix B).

2.3 Attack implementation and analysis

We selected both generalist (GLMs) and specialist (BLMs) models for benchmarking the sampling-based entity substitution attack, which requires only API access to the LLM. The selected models contain diverse encoding schemes and model architectures. Some of them have appeared on per-

[2]The limit $n \to -\infty$ indicates replacing the distractor d with k at all times. Since k is not in the perturbation set $\mathcal{E}_{\text{-}k}$, this limit is not attained.

[3]NoPtb = unperturbed (baseline), RdPtb = perturbed by random sampling, DistPtb = perturbed by distance-weighted sampling. The "-" and "+" signs indicate the two regimes ($n \ll 0$ and $n \gg 0$).

[4]https://drugcentral.org/download

[5]https://ctdbase.org/downloads/

Table 2: Adversarial testing of model performance on MedMCQA dataset

	Drug-mention questions			Disease-mention questions		
Model (size)	NoPtb	RdPtb	DistPtb (-/+)	NoPtb	RdPtb	DistPtb (-/+)
Flan-T5-xl (3B)	0.265	0.227	0.213 / 0.225	0.261	0.199	0.194 / 0.178
Flan-T5-xxl (11B)	0.291	0.254	0.234 / 0.244	0.376	0.319	0.276 / 0.286
Flan-UL2 (20B)	0.416	0.358	0.353 / 0.363	0.408	0.344	0.304 / 0.336
MedAlpaca-7B	0.433	0.383	0.381 / 0.381	0.423	0.377	0.357 / 0.374
Palmyra-Med-20B	0.675	0.602	0.585 / 0.595	0.661	0.574	0.505 / 0.548

formance leaderboards of related tasks. For GLMs, we used instruction-finetuned T5 (Flan-T5) series of models (Chung et al., 2022) and UL2 model (Tay et al., 2022) (Flan-UL2), OpenAI's GPT-3.5 (text-davinci-003) (Brown et al., 2020) and Google's PaLM2 (text-bison-001) (Anil et al., 2023). For BLMs, we used MedAlpaca-7B (Han et al., 2023) and Writer's Palmyra-Med-20B (Kamble and Alshikh, 2023). All models were evaluated at zero temperature or non-sampling setting. The prompts were structured as in Appendix C.

Implementing the framework in 2.1 for drug and disease names requires accounting for the entity types they are associated with. We used the entity type annotations according to the UMLS (Bodenreider, 2004) of the semantic groups *Chemicals & Drugs* and *Disorders*. For each question the model answers correctly, we selected the entity of the same semantic group in the distractors closest to the key as the attack target. The choice of the closest distractor is to maintain consistency and is not uniquely associated with the effects observed here. Then, we generated the replacement entity according to random or distance-weighted sampling, using the corresponding perturbation set (constructed with `FDA-drugs` or `CTD-diseases` datasets), depending on the mention label obtained in data preprocessing (see Appendix B). For each question to attack, the probability weights in Eq. (3) were calculated using the cosine distance as h between entities in a surrogate word embedding of a BERT model finetuned on biomedical text (Yuan et al., 2022),

$$h(k, e_j) = 1 - \frac{\text{Embed}(k) \cdot \text{Embed}(e_j)}{\|\text{Embed}(k)\| \cdot \|\text{Embed}(e_j)\|}. \quad (5)$$

Using a shared embedding allows model evaluation by the same standard and simplifies the identification of transferable adversarial entities between models, as will be discussed shortly. When tuning the hyperparameter n, we limited ourselves to the interval of [-50, 50] because in our experience. Beyond this region, the performance of the sampling method doesn't deviate much from that at the interval boundaries. The benchmarking results for entity substitution attacks in different LLMs are listed in Tables 1 and 2 with examples given in Appendix D. Our analysis of the attack statistics across language models and probability weights finds the following:

Two regimes of adversarial entities. The embedding space picture of the substitution attack in Fig. 1a provides a geometric interpretation of distance-weighted sampling to examine different types of adversarial examples by tuning the exponent n. We uncover two regimes, near to ($n \ll 0$) and far from ($n \gg 0$) k semantically, which correspond to the behaviors seen in Fig. 2. The difference between them is that the substituted entities are more effective (or universal) adversarial distractors than others. We use the Gini-Simpson index (Rao, 1982) to quantify the diversity of adversarial examples,

$$\eta_{\text{GS}} = 1 - \sum_j q_j^2. \quad (6)$$

In the present context, $q_j \in (0, 1]$ is the probability of occurrence of the entity e_j among all instances of adversarial entities obtained from sampling, and $\sum_j q_j = 1$. The diversity index indicates the reusability of adversarial entities and therefore their effectiveness. Fig. 2 shows that a larger positive n in Eq. (3) samples more universal adversarial entities than in other situations, leading to a marked drop in the diversity index. We call this behavior the two-regime effect and it is observed in almost all models benchmarked on both QA datasets in varying magnitude.

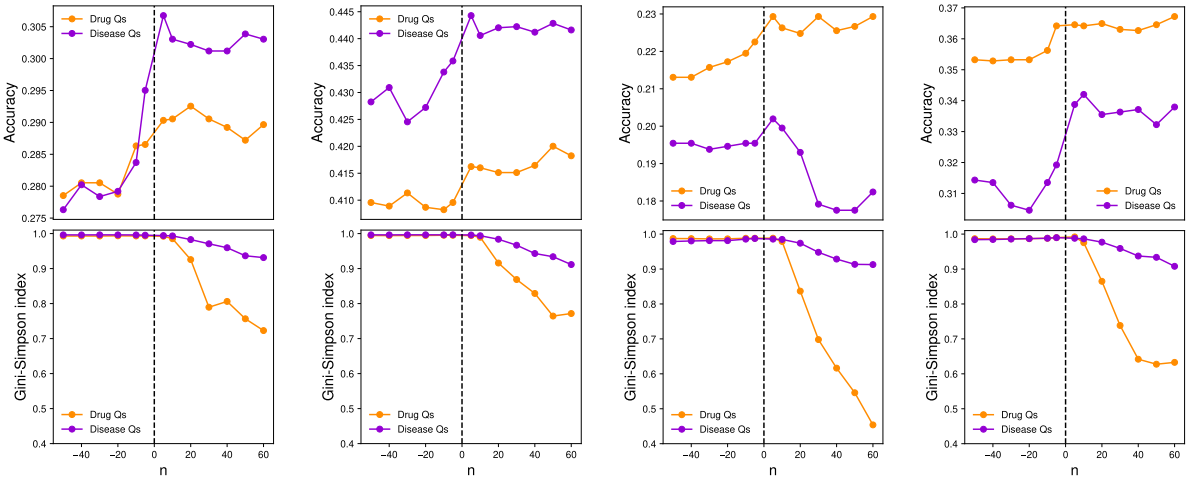

(a) (b) (c) (d)

Figure 2: Distance-weighted sampling of adversarial distractors exhibits two-regime effect at negative and positive n values (see Eq. (3)) in model accuracy and diversity index of successful attacks. The similar behaviors are compared across models and datasets in (a) Flan-T5-XXL on MedQA-USMLE, (b) MedAlpaca-7B on MedQA-USMLE, (c) Flan-T5-xl on MedMCQA, (d) Flan-UL2 on MedMCQA. The vertical dashed line indicates the equivalent location of random sampling.

The two-regime effect is the reason that sampling with probability weights either proportional (positive n) or inversely proportional (negative n) to distance h produces adversarial entities more effectively than random sampling, as shown in Tables 1 and 2. The benchmarks show that the $n \ll 0$ regime tends to have a more detrimental influence on model performance than the $n \gg 0$ regime. Effectively, at small distances, the success of the attack is related to semantic proximity (Mozes et al., 2021). Whereas at larger distances, the attack succeeds potentially due to the obscurity of the sampled entities in the question context, which may produce a profound impact on the model performance (Li et al., 2023). The reduction in the diversity index reflects the dominance of a few entities in successful attacks.

Among the adversarial entities gathered by sampling, the most effective ones are n-acetylglucosamine (an amino sugar and anti-inflammatory drug)[1] and technetium Tc 99m exametazime (a radiopharmaceutical and contrast agent)[2] for drug names, and some of the chromosome deletion syndromes and rare diseases for disease names. These entities dominate the sampling results in the $n \gg 0$ regime for all LLMs, indicating their transferability and universality. An example attack using n-acetylglucosamine as recommended medication is in Fig. 3.

[1] n-acetylglucosamine on RxList

[2] technetium Tc 99m exametazime on RxList

Entity substitution attacks manipulate explainability. Post-hoc analysis is a common way to understand the rationale behind model decisions (Murdoch et al., 2019). The Shapley values are commonly used to construct post-hoc explanation of model prediction through feature attribution (Chen et al., 2023). Here, we use *Shapley profile* to refer to the token-wise Shapley values calculated over the entire question $\mathcal{Q}$. To evaluate the effects of adversarial distractors on model explanation, we investigated the changes in the Shapley profile in the two regimes of adversarial entities for the open-source LLMs. We found that the token with the largest Shapley value (top-1 feature) is highly correlated with the model's prediction before and after adversarial perturbation, with an example shown in Fig. 3 produced using Flan-T5-large. This behavior is an indicator of the symbol-binding characteristic of LLMs in multiple-choice QA (Robinson and Wingate, 2023). Moreover, the Shapley profile typically exhibits pronounced changes before and after entity perturbation, especially in the regions close to the perturbed entities. We illustrate the point in Fig. 3 with more examples in Appendix D. Together, the observations indicate that the substitution attacks effectively manipulate model explanations, because the explanation after attack aligns with the wrong prediction, as if they were correct. The behavior occurs despite unchanged context, stem, and key ($\mathcal{C}$, $\mathcal{S}$, and k), which should contain the most important information for understanding

A 73-year-old man presents to the outpatient clinic complaining of chest pain with exertion. He states that resting for a few minutes usually resolves the chest pain. Currently, he takes 81 mg of aspirin daily. He has a blood pressure of 127/85 mm Hg and heart rate of 75/min. Physical examination reveals regular heart sounds and clear lung sounds bilateral. Which medication regimen below should be added?

A: Amlodipine daily. Sublingual nitroglycerin as needed.	A: Amlodipine daily. Sublingual nitroglycerin as needed.
B: Metoprolol and a statin daily. Sublingual nitroglycerin as needed. ✓	B: Metoprolol and a statin daily. Sublingual nitroglycerin as needed.
C: Metoprolol and ranolazine daily. Sublingual nitroglycerin as needed.	C: N-acetylglucosamine and ranolazine daily. Sublingual nitroglycerin as needed.
D: Amlodipine and a statin daily. Sublingual nitroglycerin as needed.	D: Amlodipine and a statin daily. Sublingual nitroglycerin as needed. ✓

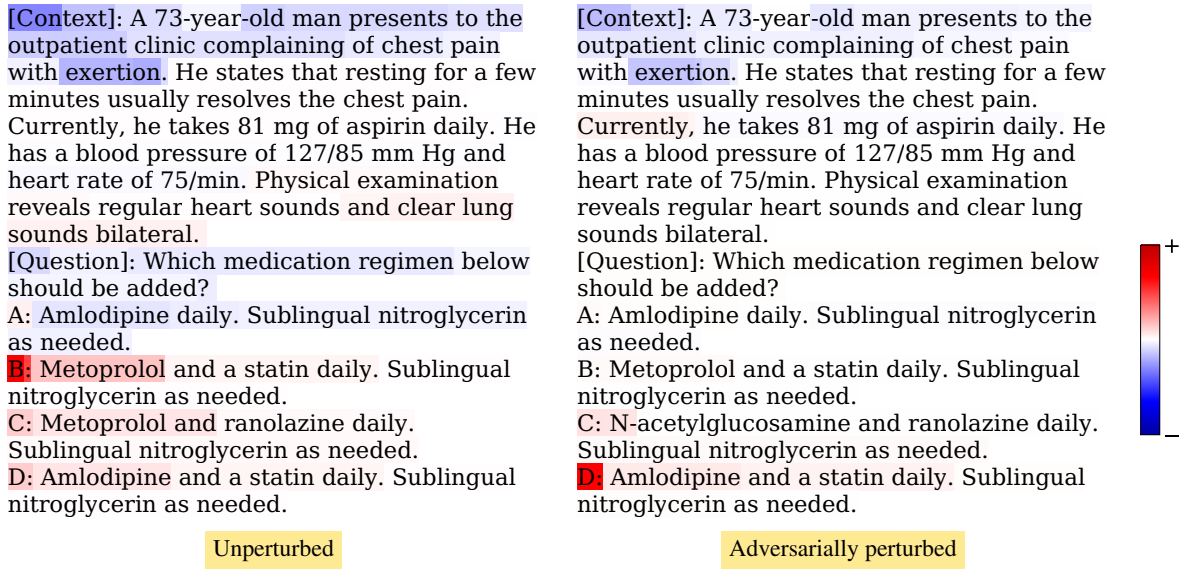

Figure 3: (Top) Example of an entity subsitution attack using n-acetylglucosamine as the adversarial distractor entity to substitute metoprolol. (Bottom) Heatmaps of token-wise Shapley values for a question before and after the adversarial attack on option C. The model prediction changes from B (unperturbed) to D (adversarially perturbed).

and answering the question.

Size is not all for model robustness. A typical metric for model robustness assessment is the change in accuracy before and after perturbation (Zhang et al., 2022). To enable comparison between datasets, we quantify the model robustness through its relative change in accuracy (Acc) upon perturbation (R-CAP), as shown partly in Fig. 4.

$$\frac{\Delta\text{Acc}}{\text{Acc}} = \frac{\Delta(n_{\text{TP}} + n_{\text{TN}})}{n_{\text{TP}} + n_{\text{TN}}} \quad (7)$$

R-CAP is equivalently the relative change in the number of correct model predictions (numbers of true positives n_{TP} and true negatives n_{TN}) (Wang et al., 2023b). For the Flan-T5 series (Chung et al., 2022), R-CAP is more pronounced for smaller and less performant models, indicating the improvement of robustness as models scale up. However, the BLMs are generally more robust to entity substitution than GLMs for models of similar size. The trend is consistent in both of the two regimes found through distance-weighted sampling. It is worth noting that although Palmyra-Med-20B (Kamble and Alshikh, 2023) has about three times the number of parameters than MedAlpaca-7B (Han et al., 2023), it is much more sensitive to entity perturbation. Therefore, the smaller BLM is more robust against perturbation. Another observation is that the R-CAP for models upon adversarial perturbations is asymmetrical. For most of the models, sampling adversarial entities closer to the key induces more performance drop than sampling fur-

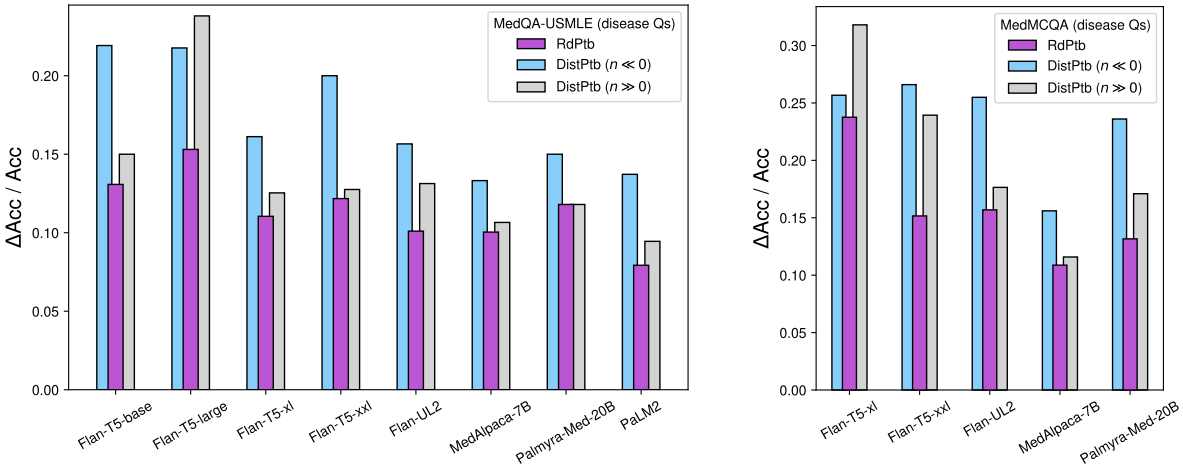

(a) (b)

Figure 4: Model robustness evaluation on disease-mention questions from (a) MedQA-USMLE and (b) MedMCQA datasets. The horizontal order partially reflects model performance without perturbation.

ther away.

3 Related works

Text substitution attack Adversarial attacks in neural language models have been demonstrated from the character to the sentence level (Zhang et al., 2020). The majority of text substitution attacks employ typos (Gao et al., 2018) and synonyms (Mozes et al., 2021), which effectively impose lexical constraints. Word frequency has been used as a metric to select adversarial examples in (Mozes et al., 2021). Synonym substitution attacks have been demonstrated in sentiment analysis (Liu et al., 2022). (Alzantot et al., 2018) used a genetic algorithm to search for adversarial word substitution. Model sensitivity to entity-level perturbations such as substitution and swapping create problematic situations for knowledge-intensive tasks, such as knowledge conflicts (Longpre et al., 2021), compromised performance in machine reading comprehension (Yan et al., 2022) and table interpretation tasks (Koleva et al., 2023).

Robustness in QA The robustness of QA systems has been studied using perturbative methods (Jia and Liang, 2017) and meaning-preserving transformations (Gan and Ng, 2019; Elazar et al., 2021). (Sen and Saffari, 2020) introduced a set of perturbations to assess language model generalization between QA datasets. (Richardson and Sabharwal, 2020) proposed a distractor design method using proximity in semantic space to probe model understanding of word senses. (Awadalla et al., 2022) examined the connection between in- and out-of-distribution robustness for different LLMs. The performance of LLMs is sensitive to the ordering of options (Pezeshkpour and Hruschka, 2023; Zheng et al., 2023). Overall, the brittleness of knowledge in LLMs remains a primary limitation for their applications (Razniewski et al., 2021; Elazar et al., 2021; Ji et al., 2023).

Adversarial attack characteristics The defining characteristic of adversarial attacks is to induce false predictions in trained models. Besides, attacks may possess transferability, universality (Zou et al., 2023), and they may obstruct model explanation (Noppel and Wressnegger, 2024). Transferability refers to the effectiveness of the same example being adversarial to a number of distinct models (Papernot et al., 2016). Universal adversarial perturbations are *reusable* across many examples tested on the same model (Moosavi-Dezfooli et al., 2017; Wallace et al., 2019a). Model explanations may be manipulated by specifically designed perturbations to the inputs (Dombrowski et al., 2019) or biased data subsampling (Laberge et al., 2023).

4 Discussion and future work

The present study illustrates that side information introduced through entity substitution leads to noticeable degradation of model performance. Although perturbations at the distractor level have fewer restrictions than other parts of the question, we maintain the entity type consistency within the distractor to generate natural adversarial perturbations. The setting simulates realistic scenarios where the “attacks" may be initiated by unsuspect-

ing users, such as biomedical professionals using an LLM-based clinical decision support system (Liu et al., 2023). In the examples illustrated in Fig. 3 and Appendix D, adversarial distractors can lead to disease misdiagnosis or prescription of the wrong medication, which are detrimental scenarios that may occur in algorithmic decision-making in biomedicine. Our approach may be integrated into interactive platforms for adversarial testing in open-ended human-machine conversation (Chao et al., 2023) for online model auditing.

Future research may focus on how to make LLMs more reliable through prompt engineering (Si et al., 2023; Nori et al., 2023) or leveraging retrieval from nonparametric knowledge bases (Weikum et al., 2021; Soman et al., 2023) or text sources (Wang et al., 2023c) to improve model generalization to long-tail entities. The distance-weighted sampling may be extended to kernel methods (Hofmann et al., 2008) or metric learning (Bellet et al., 2015), which can capture more general distance dependence. Besides vulnerability identification, the current work motivates adversarial defense strategies based on generalized distances to establish a tiered system for validating adversarial robustness in user-facing and domain-oriented applications.

Limitations

Although NEs convey parts of the key information in the question, they fall short of capturing the entirety of the textual meaning. To explore other attack surfaces within the adversarial landscape of QA, one would need to evaluate the robustness in multidimensional perturbation space (Wang et al., 2022). Besides, the meaning of Shapley values calculated here requires additional interpretation due to feature correlation. Although we use them in part to reveal the importance of specific tokens, their robustness is also worth exploring using sentence-level features (Chen et al., 2023).

Ethics statement

The scenarios considered in this work using public datasets relate to potential real-world deployment of LLMs in biomedical education and clinical decision support. Our results show that the lack of adversarial testing methods in high-stakes domains may incur serious consequences because of naturally looking adversarial examples. The design and deployment of adversarial testing should be specific to and carried out at every stage of the model development cycle, paired with tools such as adversarially robust model explainability methods for understanding model behavior.

References

2023. Will ChatGPT transform healthcare? *Nature Medicine*, 29(3):505–506.

Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. Generating Natural Language Adversarial Examples. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Stroudsburg, PA, USA. Association for Computational Linguistics.

Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. 2023. Palm 2 technical report.

Anas Awadalla, Mitchell Wortsman, Gabriel Ilharco, Sewon Min, Ian Magnusson, Hannaneh Hajishirzi, and Ludwig Schmidt. 2022. Exploring The Landscape of Distributional Robustness for Question Answering

Models. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5971–5987, Stroudsburg, PA, USA. Association for Computational Linguistics.

Aurélien Bellet, Amaury Habrard, and Marc Sebban. 2015. *Metric Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Springer International Publishing, Cham.

O. Bodenreider. 2004. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Research*, 32(90001):267D–270.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Dhawaleswar Rao CH and Sujan Kumar Saha. 2020. Automatic Multiple Choice Question Generation From Text: A Survey. *IEEE Transactions on Learning Technologies*, 13(1):14–25.

Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J. Pappas, and Eric Wong. 2023. Jailbreaking Black Box Large Language Models in Twenty Queries. *arXiv*, page 2310.08419.

Hugh Chen, Ian C. Covert, Scott M. Lundberg, and Su-In Lee. 2023. Algorithms to estimate Shapley value feature attributions. *Nature Machine Intelligence*, 5(6):590–601.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. Scaling instruction-finetuned language models.

Peter Clark and Oren Etzioni. 2016. My Computer Is an Honor Student — But How Intelligent Is It? Standardized Tests as a Measure of AI. *AI Magazine*, 37(1):5–12.

A. P. Davis, C. G. Murphy, C. A. Saraceni-Richards, M. C. Rosenstein, T. C. Wiegers, and C. J. Mattingly. 2009. Comparative Toxicogenomics Database: a knowledgebase and discovery tool for chemical-gene-disease networks. *Nucleic Acids Research*, 37(Database):D786–D792.

Ann-Kathrin Dombrowski, Maximillian Alber, Christopher Anders, Marcel Ackermann, Klaus-Robert Müller, and Pan Kessel. 2019. Explanations can be manipulated and geometry is to blame. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard Hovy, Hinrich Schütze, and Yoav Goldberg. 2021. Measuring and Improving Consistency in Pretrained Language Models. *Transactions of the Association for Computational Linguistics*, 9:1012–1031.

Wee Chung Gan and Hwee Tou Ng. 2019. Improving the Robustness of Question Answering Systems to Question Paraphrasing. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6065–6075, Stroudsburg, PA, USA. Association for Computational Linguistics.

Ji Gao, Jack Lanchantin, Mary Lou Soffa, and Yanjun Qi. 2018. Black-Box Generation of Adversarial Text Sequences to Evade Deep Learning Classifiers. In *2018 IEEE Security and Privacy Workshops (SPW)*, pages 50–56. IEEE.

Matt Gardner, Jonathan Berant, Hannaneh Hajishirzi, Alon Talmor, and Sewon Min. 2019. Question Answering is a Format; When is it Useful? *arXiv*, page 1909.11291.

Tianyu Han, Lisa C. Adams, Jens-Michalis Papaioannou, Paul Grundmann, Tom Oberhauser, Alexander Löser, Daniel Truhn, and Keno K. Bressem. 2023. Medalpaca – an open-source collection of medical conversational ai models and training data.

Thomas Hofmann, Bernhard Schölkopf, and Alexander J. Smola. 2008. Kernel methods in machine learning. *The Annals of Statistics*, 36(3):1171–1220.

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys*, 55(12):1–38.

Robin Jia and Percy Liang. 2017. Adversarial Examples for Evaluating Reading Comprehension Systems. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Stroudsburg, PA, USA. Association for Computational Linguistics.

Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What Disease Does This Patient Have? A Large-Scale Open Domain Question Answering Dataset from Medical Exams. *Applied Sciences*, 11(14):6421.

Kiran Kamble and Waseem Alshikh. 2023. Palmyra-med: Instruction-based fine-tuning of llms enhancing medical domain performance.

Aneta Koleva, Martin Ringsquandl, and Volker Tresp. 2023. Adversarial Attacks on Tables with Entity Swap. In *Joint Proceedings of Workshops at the 49th International Conference on Very Large Data Bases (VLDB 2023), Vancouver, Canada*, volume 3462 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Tiffany H. Kung, Morgan Cheatham, Arielle Medenilla, Czarina Sillos, Lorie De Leon, Camille Elepaño, Maria Madriaga, Rimel Aggabao, Giezel Diaz-Candido, James Maningo, and Victor Tseng. 2023. Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. *PLOS Digital Health*, 2(2):e0000198.

Gabriel Laberge, Ulrich Aïvodji, Satoshi Hara, Mario Marchand, and Foutse Khomh. 2023. Fooling SHAP with Stealthily Biased Sampling. In *The Eleventh International Conference on Learning Representations*.

Eric Lehman, Evan Hernandez, Diwakar Mahajan, Jonas Wulff, Micah J Smith, Zachary Ziegler, Daniel Nadler, Peter Szolovits, Alistair Johnson, and Emily Alsentzer. 2023. Do We Still Need Clinical Language Models? In *Proceedings of the Conference on Health, Inference, and Learning*, volume 209 of *Proceedings of Machine Learning Research*, pages 578–597. PMLR.

Xinzhe Li, Ming Liu, Shang Gao, and Wray Buntine. 2023. A Survey on Out-of-Distribution Evaluation of Neural NLP Models. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 6683–6691, California. International Joint Conferences on Artificial Intelligence Organization.

Aiwei Liu, Honghai Yu, Xuming Hu, Shu'ang Li, Li Lin, Fukun Ma, Yawen Yang, and Lijie Wen. 2022. Character-level White-Box Adversarial Attacks against Transformers via Attachable Subwords Substitution. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7664–7676, Stroudsburg, PA, USA. Association for Computational Linguistics.

Siru Liu, Aileen P Wright, Barron L Patterson, Jonathan P Wanderer, Robert W Turer, Scott D Nelson, Allison B McCoy, Dean F Sittig, and Adam Wright. 2023. Assessing the Value of ChatGPT for Clinical Decision Support Optimization. *medRxiv*.

Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. Entity-Based Knowledge Conflicts in Question Answering. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063, Stroudsburg, PA, USA. Association for Computational Linguistics.

Danaë Metaxa, Joon Sung Park, Ronald E. Robertson, Karrie Karahalios, Christo Wilson, Jeff Hancock, and Christian Sandvig. 2021. Auditing Algorithms: Understanding Algorithmic Systems from the Outside In. *Foundations and Trends® in Human–Computer Interaction*, 14(4):272–344.

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. 2017. Universal Adversarial Perturbations. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 86–94. IEEE.

Maximilian Mozes, Pontus Stenetorp, Bennett Kleinberg, and Lewis Griffin. 2021. Frequency-Guided Word Substitutions for Detecting Textual Adversarial Examples. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 171–186, Stroudsburg, PA, USA. Association for Computational Linguistics.

W. James Murdoch, Chandan Singh, Karl Kumbier, Reza Abbasi-Asl, and Bin Yu. 2019. Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences*, 116(44):22071–22080.

Mark Neumann, Daniel King, Iz Beltagy, and Waleed Ammar. 2019. ScispaCy: Fast and Robust Models for Biomedical Natural Language Processing. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 319–327, Stroudsburg, PA, USA. Association for Computational Linguistics.

Maximilian Noppel and Christian Wressnegger. 2024. SoK: Explainable Machine Learning in Adversarial Environments. In *2024 IEEE Symposium on Security and Privacy (SP)*, page 21. IEEE Computer Society.

Harsha Nori, Yin Tat Lee, Sheng Zhang, Dean Carignan, Richard Edgar, Nicolo Fusi, Nicholas King, Jonathan Larson, Yuanzhi Li, Weishung Liu, Renqian Luo, Scott Mayer McKinney, Robert Osazuwa Ness, Hoifung Poon, Tao Qin, Naoto Usuyama, Chris White, and Eric Horvitz. 2023. Can Generalist Foundation Models Outcompete Special-Purpose Tuning? Case Study in Medicine. *arXiv*, page 2311.16452.

Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. MedMCQA: A Large-scale Multi-Subject Multi-Choice Dataset for Medical domain Question Answering. In *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 248–260. PMLR.

Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. 2016. Transferability in Machine Learning: from Phenomena to Black-Box Attacks using Adversarial Samples. *arXiv*, page 1605.07277.

Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and

Alexander Miller. 2019. Language Models as Knowledge Bases? 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 2463–2473, Stroudsburg, PA, USA. Association for Computational Linguistics.

Pouya Pezeshkpour and Estevam Hruschka. 2023. Large Language Models Sensitivity to The Order of Options in Multiple-Choice Questions. *arXiv*, page 2308.11483.

C.Radhakrishna Rao. 1982. Diversity and dissimilarity coefficients: A unified approach. *Theoretical Population Biology*, 21(1):24–43.

Simon Razniewski, Andrew Yates, Nora Kassner, and Gerhard Weikum. 2021. Language Models As or For Knowledge Bases. *arXiv*, page 2110.04888.

Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. 2019. Generating Natural Language Adversarial Examples through Probability Weighted Word Saliency. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1085–1097, Stroudsburg, PA, USA. Association for Computational Linguistics.

Kyle Richardson and Ashish Sabharwal. 2020. What Does My QA Model Know? Devising Controlled Probes Using Expert Knowledge. *Transactions of the Association for Computational Linguistics*, 8:572–588.

Joshua Robinson and David Wingate. 2023. Leveraging Large Language Models for Multiple Choice Question Answering. In *The Eleventh International Conference on Learning Representations*.

Priyanka Sen and Amir Saffari. 2020. What do Models Learn from Question Answering Datasets? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2429–2438, Stroudsburg, PA, USA. Association for Computational Linguistics.

Erfan Shayegani, Md Abdullah Al Mamun, Yu Fu, Pedram Zaree, Yue Dong, and Nael Abu-Ghazaleh. 2023. Survey of Vulnerabilities in Large Language Models Revealed by Adversarial Attacks. *arXiv*, page 2310.10844.

Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Lee Boyd-Graber, and Lijuan Wang. 2023. Prompting GPT-3 To Be Reliable. In *The Eleventh International Conference on Learning Representations*.

Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Abubakr Babiker, Nathanael Schärli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, Blaise Agüera y Arcas, Dale Webster, Greg S. Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomasev, Yun Liu, Alvin Rajkomar, Joelle Barral, Christopher Semturs, Alan Karthikesalingam, and Vivek Natarajan. 2023a. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180.

Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, Mike Schaekermann, Amy Wang, Mohamed Amin, Sami Lachgar, Philip Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Aguera y Arcas, Nenad Tomasev, Yun Liu, Renee Wong, Christopher Semturs, S. Sara Mahdavi, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Shekoofeh Azizi, Alan Karthikesalingam, and Vivek Natarajan. 2023b. Towards Expert-Level Medical Question Answering with Large Language Models. *arXiv*, page 2305.09617.

Karthik Soman, Peter W Rose, John H Morris, Rabia E Akbas, Brett Smith, Braian Peetoom, Catalina Villouta-Reyes, Gabriel Cerono, Yongmei Shi, Angela Rizk-Jackson, Sharat Israni, Charlotte A Nelson, Sui Huang, and Sergio E Baranzini. 2023. Biomedical knowledge graph-enhanced prompt generation for large language models. *arXiv*, page 2311.17330.

Giovanni Spitale, Nikola Biller-Andorno, and Federico Germani. 2023. AI model GPT-3 (dis)informs us better than humans. *Science Advances*, 9(26).

Mujeen Sung, Jinhyuk Lee, Sean Yi, Minji Jeon, Sungdong Kim, and Jaewoo Kang. 2021. Can Language Models be Biomedical Knowledge Bases? In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4723–4734, Stroudsburg, PA, USA. Association for Computational Linguistics.

Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Siamak Shakeri, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Denny Zhou, Neil Houlsby, and Donald Metzler. 2022. UL2: Unifying Language Learning Paradigms. *arXiv*, page 2205.05131.

Robert Tinn, Hao Cheng, Yu Gu, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2023. Fine-tuning large neural language models for biomedical natural language processing. *Patterns*, 4(4):100729.

Oleg Ursu, Jayme Holmes, Jeffrey Knockel, Cristian G. Bologa, Jeremy J. Yang, Stephen L. Mathias, Stuart J. Nelson, and Tudor I. Oprea. 2017. DrugCentral: online drug compendium. *Nucleic Acids Research*, 45(D1):D932–D939.

Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019a. Universal Adversarial Triggers for Attacking and Analyzing NLP. 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 2153–2162, Stroudsburg, PA, USA. Association for Computational Linguistics.

Eric Wallace, Pedro Rodriguez, Shi Feng, Ikuya Yamada, and Jordan Boyd-Graber. 2019b. Trick Me If You Can: Human-in-the-Loop Generation of Adversarial Examples for Question Answering. *Transactions of the Association for Computational Linguistics*, 7:387–401.

Benyou Wang, Qianqian Xie, Jiahuan Pei, Zhihong Chen, Prayag Tiwari, Zhao Li, and Jie Fu. 2023a. Pre-trained Language Models in Biomedical Domain: A Systematic Survey. *ACM Computing Surveys*.

Boxin Wang, Chejian Xu, Xiangyu Liu, Yu Cheng, and Bo Li. 2022. SemAttack: Natural Textual Attacks via Different Semantic Spaces. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 176–205, Stroudsburg, PA, USA. Association for Computational Linguistics.

Haoyu Wang, Guozheng Ma, Cong Yu, Ning Gui, Linrui Zhang, Zhiqi Huang, Suwei Ma, Yongzhe Chang, Sen Zhang, Li Shen, Xueqian Wang, Peilin Zhao, and Dacheng Tao. 2023b. Are Large Language Models Really Robust to Word-Level Perturbations? *arXiv*, page 2309.11166.

Yubo Wang, Xueguang Ma, and Wenhu Chen. 2023c. Augmenting Black-box LLMs with Medical Textbooks for Clinical Question Answering. *arXiv*, page 2309.02233.

Chih-Hsuan Wei, Alexis Allot, Robert Leaman, and Zhiyong Lu. 2019. PubTator central: automated concept annotation for biomedical full text articles. *Nucleic Acids Research*, 47(W1):W587–W593.

Gerhard Weikum, Xin Luna Dong, Simon Razniewski, and Fabian Suchanek. 2021. Machine Knowledge: Creation and Curation of Comprehensive Knowledge Bases. *Foundations and Trends® in Databases*, 10(2-4):108–490.

Michael Wornow, Yizhe Xu, Rahul Thapa, Birju Patel, Ethan Steinberg, Scott Fleming, Michael A. Pfeffer, Jason Fries, and Nigam H. Shah. 2023. The shaky foundations of large language models and foundation models for electronic health records. *npj Digital Medicine*, 6(1):135.

Jun Yan, Yang Xiao, Sagnik Mukherjee, Bill Yuchen Lin, Robin Jia, and Xiang Ren. 2022. On the Robustness of Reading Comprehension Models to Entity Renaming. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 508–520, Stroudsburg, PA, USA. Association for Computational Linguistics.

Puyudi Yang, Jianbo Chen, Cho-Jui Hsieh, Jane-Ling Wang, and Michael I Jordan. 2020. Greedy Attack and Gumbel Attack: Generating Adversarial Examples for Discrete Data. *Journal of Machine Learning Research*, 21(43):1–36.

Zheng Yuan, Zhengyun Zhao, Haixia Sun, Jiao Li, Fei Wang, and Sheng Yu. 2022. CODER: Knowledge-infused cross-lingual medical term embedding for term normalization. *Journal of Biomedical Informatics*, 126:103983.

Jie M. Zhang, Mark Harman, Lei Ma, and Yang Liu. 2022. Machine Learning Testing: Survey, Landscapes and Horizons. *IEEE Transactions on Software Engineering*, 48(1):1–36.

Wei Emma Zhang, Quan Z. Sheng, Ahoud Alhazmi, and Chenliang Li. 2020. Adversarial Attacks on Deep-learning Models in Natural Language Processing. *ACM Transactions on Intelligent Systems and Technology*, 11(3):1–41.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. A Survey of Large Language Models. *arXiv*, page 2303.18223.

Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. Large Language Models Are Not Robust Multiple Choice Selectors. *arXiv*, page 2309.03882v3.

Andy Zou, Zifan Wang, J. Zico Kolter, and Matt Fredrikson. 2023. Universal and Transferable Adversarial Attacks on Aligned Language Models. *arXiv*, page 2307.15043.

A Proof of the sampling regimes

Here we present a simple proof of the deterministic limit of the distance-weighted sampling described in Eq. (3) of the main text. Although the distance function h is not given a specific form, we assume it to satisfy the basic properties of a metric. Given a perturbation set $\mathcal{E}_{\text{-}k} = \{e_1, ..., e_m\}$ with respect to a key k, there exists a corresponding set of distance values $\mathcal{H}_{\text{-}k}$.

A.1 The limit of $n \to +\infty$

We order the elements of $\mathcal{H}_{\text{-}k}$ and define the ordered set $\mathcal{H}_{\text{-}k,\uparrow} = \{h_{1,\uparrow}, ..., h_{m,\uparrow}\}$ such that,

$$h_{i,\uparrow} < h_{j,\uparrow} \quad \text{iff} \quad i < j \tag{8}$$

This indicates that $h_{m,\uparrow} = \max \mathcal{H}_{\text{-}k,\uparrow}$, and $h_{j,\uparrow}/h_{m,\uparrow} < 1$ for any $j \neq m$. We compute the probability weights as $n \to +\infty$ for each entity $e_j \in \mathcal{E}_{\text{-}k}$. When $j \neq m$,

$$\begin{aligned}
&\lim_{n\to+\infty} p_{j,\uparrow} = \lim_{n\to+\infty} \frac{h_{j,\uparrow}^n}{\sum_j h_{j,\uparrow}^n} \\
&= \lim_{n\to+\infty} \frac{(h_{j,\uparrow}/h_{m,\uparrow})^n}{\sum_{j\neq m}(h_{j,\uparrow}/h_{m,\uparrow})^n + (h_{m,\uparrow}/h_{m,\uparrow})^n} \\
&= \frac{0^n}{0^n + 1^n} = 0.
\end{aligned} \tag{9}$$

Similarly, as $n \to +\infty$, the probability weight for the mth entity is

$$\lim_{n\to+\infty} p_{m,\uparrow} = 1. \tag{10}$$

For the entities in the perturbation set $\mathcal{E}_{\text{-}k}$, this indicates that only

$$e_j = \operatorname{argmax} \mathcal{H}_{\text{-}k,\uparrow} = \operatorname*{argmax}_{e_j \in \mathcal{E}_{\text{-}k}} h(k, e_j) \tag{11}$$

is selected with certainty at the deterministic limit.

A.2 The limit of $n \to -\infty$

A similar approach to B.1 can be used to construct the probability weights in the limit $n \to -\infty$. As noted in the footnote of the main text, this limit is not attained since $k \notin \mathcal{E}_{\text{-}k}$.

A.3 The case of $n = 0$

When $n = 0$, it is straightforward to show that $p_j|_{n=0} = 1/m$, regardless of j using the definition in Eq. (3). This scenario therefore corresponds to uniform random sampling.

B Biomedical entity processing

The original `MedQA-USMLE` (Jin et al., 2021) and `MedMCQA` (Pal et al., 2022) datasets lack entity annotations. In the present work, we carried out a two-step annotation for both datasets.

UMLS entity types We annotated the datasets using the type unique identifiers (TUIs) specified in the UMLS (Bodenreider, 2004) for biomedical entities. The annotation produces TUIs and the entity boundaries for entire questions and was carried out using `scispaCy` (Neumann et al., 2019). The annotation should be regarded as algorithmically produced, rather than human-annotated and proofread ground truth.

Binary mention labels For entity-level perturbation, the distinction between questions with and disease entity mentions is based on the type of entities present in the distractors. We generate the binary mention labels ("Drugs?" and "Diseases?") for each question selected from the two datasets using the UMLS semantic groups of the TUIs in (i) for the distractors. Specifically, a drug-mention question has one or more entities in the distractors of the *Chemicals & Drugs* semantic group. A disease-mention question has one or more entities of the *Disorders* semantic group in its distractors.

C Prompt formatting

We structured the prompt with bracketed tags to divide different parts (Context, Question, Answer). An example taken from the MedQA-USMLE dataset (Jin et al., 2021) is shown below. For the MedMCQA dataset (Pal et al., 2022), only the Context and Answer[1] tags were used because the questions are short and generally contain no context. The prompt formatting improves model performance and facilitates answer parsing in the model-generated text. The increased space between paragraphs in the following is only to assist viewing and is not included in the actual prompt.

D Model evaluation and attack examples

We list some examples of successful entity substitution attacks using adversarial distractors following Fig. 3 in the main text. For each example, the text box on top (the prompt formatting removed) shows the original options on the left side and the

[1] Using the Context tag improves model performance than the Question tag here.

Prompt format

Answer the following question without explanation.

[Content]: A 23-year-old pregnant woman at 22 weeks gestation presents with burning upon urination. She states it started 1 day ago and has been worsening despite drinking more water and taking cranberry extract. She otherwise feels well and is followed by a doctor for her pregnancy. Her temperature is 97.7°F (36.5°C), blood pressure is 122/77 mmHg, pulse is 80/min, respirations are 19/min, and oxygen saturation is 98% on room air. Physical exam is notable for an absence of costovertebral angle tenderness and a gravid uterus.

[Question]: Which of the following are the best treatment for this patient?
A: Ampicillin
B: Ceftriaxone
C: Doxycycline
D: Nitrofurantoin

[Answer]:

perturbed options on the right side. The entity of interest before and after perturbation is colored in red and blue, respectively. The checkmark indicates the model prediction on unperturbed and perturbed questions. Beneath the text box are the corresponding heatmaps of token-wise Shapley values (or Shapley profile) before and after the entity substitution attack. The Shapley profiles for the MedQA-USMLE dataset (Jin et al., 2021) were generated using the model Flan-T5-large and those for the MedMCQA (Pal et al., 2022) using the model Flan-T5-xxl (Chung et al., 2022). The Shapley values were calculated using the partition explainer in `shap`, which uses the Owen values as an approximation (Chen et al., 2023).

Of the LLMs evaluated in this work, the proprietary models (GPT-3.5, PaLM2) were accessed through third-party APIs, while the open-source models were accessed on local cache to improve speed. The evaluations were carried out on a variety of NVIDIA GPUs depending on the model size. For smaller model such as the Flan-T5 series (Chung et al., 2022), the runtime on each type of questions (involving drug or disease mentions) in a dataset was generally between 10 and 30 minutes, while it was from one to more than four hours for larger model such Flan-UL2 (Tay et al., 2022) and Palmyra-Med-20B (Kamble and Alshikh, 2023). For proprietary models accessed through their APIs, the runtime for each type of questions is between two and four hours.

MedQA-USMLE

An 18-year-old man presents with a sudden loss of consciousness while playing college football. There was no history of a concussion. Echocardiography shows left ventricular hypertrophy and increased thickness of the interventricular septum. Which is the most likely pathology underlying the present condition?

A: Mutation in the myosin heavy chain ✓	A: Mutation in the myosin heavy chain
B: Drug abuse	B: Drug abuse
C: Viral infection	C: Muller barth menger syndrome ✓
D: Autoimmunity of myocardial fibers	D: Autoimmunity of myocardial fibers

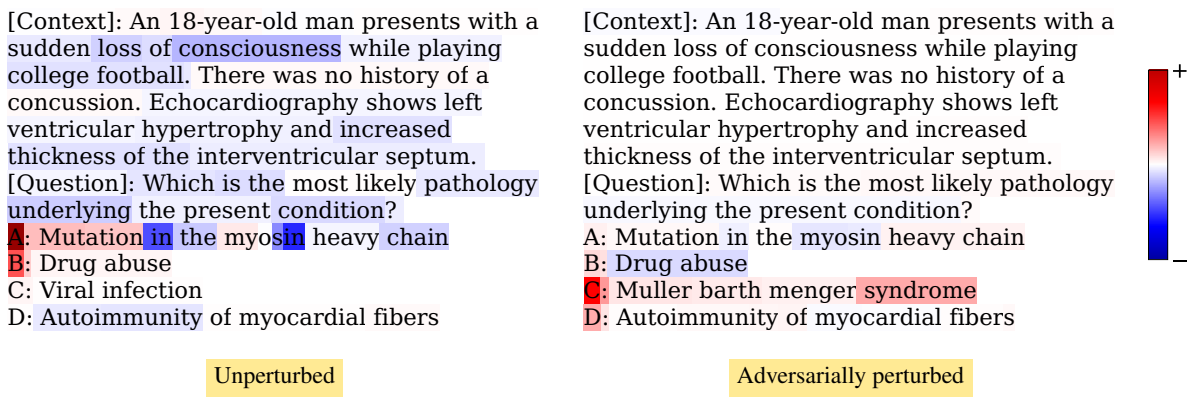

MedQA-USMLE

A 41-year-old woman comes to the physician because of a 3-month history of anxiety, difficulty falling asleep, heat intolerance, and a 6-kg (13.2-lb) weight loss. The patient's nephew, who is studying medicine, mentioned that her symptoms might be caused by a condition that is due to somatic activating mutations of the genes for the TSH receptor. Examination shows warm, moist skin and a 2-cm, nontender, subcutaneous mass on the anterior neck. Which of the following additional findings should most raise concern for a different underlying etiology of her symptoms?

A: Nonpitting edema ✓	A: Nonpitting edema
B: Atrial fibrillation	B: Atrial fibrillation
C: Lid lag	C: Lid lag
D: Fine tremor	D: Nonpuerperal galactorrhea ✓

[Context]: A 41-year-old woman comes to the physician because of a 3-month history of anxiety, difficulty falling asleep, heat intolerance, and a 6-kg (13.2-lb) weight loss. The patient's nephew, who is studying medicine, mentioned that her symptoms might be caused by a condition that is due to somatic activating mutations of the genes for the TSH receptor. Examination shows warm, moist skin and a 2-cm, nontender, subcutaneous mass on the anterior neck.
[Question]: Which of the following additional findings should most raise concern for a different underlying etiology of her symptoms?
A: Nonpitting edema
B: Atrial fibrillation
C: Lid lag
D: Fine tremor

Unperturbed

[Context]: A 41-year-old woman comes to the physician because of a 3-month history of anxiety, difficulty falling asleep, heat intolerance, and a 6-kg (13.2-lb) weight loss. The patient's nephew, who is studying medicine, mentioned that her symptoms might be caused by a condition that is due to somatic activating mutations of the genes for the TSH receptor. Examination shows warm, moist skin and a 2-cm, nontender, subcutaneous mass on the anterior neck.
[Question]: Which of the following additional findings should most raise concern for a different underlying etiology of her symptoms?
A: Nonpitting edema
B: Atrial fibrillation
C: Lid lag
D: Nonpuerperal galactorrhea

Adversarially perturbed

MedQA-USMLE

A 25-year-old man is brought to the emergency department after his girlfriend discovered him at home in a minimally responsive state. He has a history of drinking alcohol excessively and using illicit drugs. On arrival, he does not respond to commands but withdraws all extremities to pain. His pulse is 90/min, respirations are 8/min, and blood pressure is 130/90 mm Hg. Pulse oximetry while receiving bag-valve-mask ventilation shows an oxygen saturation of 95%. Examination shows cool, dry skin, with scattered track marks on his arms and legs. The pupils are pinpoint and react sluggishly to light. His serum blood glucose level is 80 mg/dL. The most appropriate next step in management is intravenous administration of which of the following?

A: Naloxone
B: Phentolamine ✓
C: Methadone
D: Naltrexone

A: Naloxone
B: Phentolamine
C: Methadone ✓
D: Tienilic acid

[Context]: A 25-year-old man is brought to the emergency department after his girlfriend discovered him at home in a minimally responsive state. He has a history of drinking alcohol excessively and using illicit drugs. On arrival, he does not respond to commands but withdraws all extremities to pain. His pulse is 90/min, respirations are 8/min, and blood pressure is 130/90 mm Hg. Pulse oximetry while receiving bag-valve-mask ventilation shows an oxygen saturation of 95%. Examination shows cool, dry skin, with scattered track marks on his arms and legs. The pupils are pinpoint and react sluggishly to light. His serum blood glucose level is 80 mg/dL.
[Question]: The most appropriate next step in management is intravenous administration of which of the following?
A: Naloxone
B: Phentolamine
C: Methadone
D: Naltrexone

Unperturbed

[Context]: A 25-year-old man is brought to the emergency department after his girlfriend discovered him at home in a minimally responsive state. He has a history of drinking alcohol excessively and using illicit drugs. On arrival, he does not respond to commands but withdraws all extremities to pain. His pulse is 90/min, respirations are 8/min, and blood pressure is 130/90 mm Hg. Pulse oximetry while receiving bag-valve-mask ventilation shows an oxygen saturation of 95%. Examination shows cool, dry skin, with scattered track marks on his arms and legs. The pupils are pinpoint and react sluggishly to light. His serum blood glucose level is 80 mg/dL.
[Question]: The most appropriate next step in management is intravenous administration of which of the following?
A: Naloxone
B: Phentolamine
C: Methadone
D: Tienilic acid

Adversarially perturbed

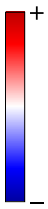

MedQA-USMLE

A 70-year-old woman presents to the office for a yearly physical. She states she has recently started experiencing pain in her legs and her back. Last year, she experienced a fracture of her left arm while trying to lift groceries. The patient states that she does not consume any dairy and does not go outside often because of the pain in her legs and back. Of note, she takes carbamazepine for seizures. On exam, her vitals are within normal limits. You suspect the patient might have osteomalacia. Testing for which of the following is the next best step to confirm your suspicion?

A: 25-hydroxyvitamin D ✓
B: 1,25-hydroxyvitamin D
C: Pre-vitamin D3
D: Dietary vitamin D2

A: 25-hydroxyvitamin D
B: 1,25-hydroxyvitamin D
C: Pre-vitamin D3 ✓
D: Stiripentol

[Context]: A 70-year-old woman presents to the office for a yearly physical. She states she has recently started experiencing pain in her legs and her back. Last year, she experienced a fracture of her left arm while trying to lift groceries. The patient states that she does not consume any dairy and does not go outside often because of the pain in her legs and back. Of note, she takes carbamazepine for seizures. On exam, her vitals are within normal limits. You suspect the patient might have osteomalacia.
[Question]: Testing for which of the following is the next best step to confirm your suspicion?
A: 25-hydroxyvitamin D
B: 1,25-hydroxyvitamin D
C: Pre-vitamin D3
D: Dietary vitamin D2

Unperturbed

[Context]: A 70-year-old woman presents to the office for a yearly physical. She states she has recently started experiencing pain in her legs and her back. Last year, she experienced a fracture of her left arm while trying to lift groceries. The patient states that she does not consume any dairy and does not go outside often because of the pain in her legs and back. Of note, she takes carbamazepine for seizures. On exam, her vitals are within normal limits. You suspect the patient might have osteomalacia.
[Question]: Testing for which of the following is the next best step to confirm your suspicion?
A: 25-hydroxyvitamin D
B: 1,25-hydroxyvitamin D
C: Pre-vitamin D3
D: Stiripentol

Adversarially perturbed

MedMCQA

Which of the following is granted orphan drug status for the treatment of Dravet's syndrome?

A: Stiripentol	A: Stiripentol
B: Icatibant	B: Tocilizumab ✓
C: Pre-vitamin D3	C: Pitolisant
D: Tafamidis ✓	D: Tafamidis

[Context]: Which of the following is granted orphan drug status for the treatment of Dravet's syndrome?
A: Stiripentol
B: Icatibant
C: Pitolisant
D: Tafamidis

Unperturbed

[Context]: Which of the following is granted orphan drug status for the treatment of Dravet's syndrome?
A: Stiripentol
B: Tocilizumab
C: Pitolisant
D: Tafamidis

Adversarially perturbed

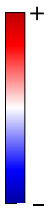

MedMCQA

Which among the following may be used as a sedatives-hypnotic?

A: Zolmitriptan	A: Zolmitriptan ✓
B: Zileuton	B: Zileuton
C: Zolpidem ✓	C: N-acetylglucosamine
D: Zalcitabine	D: Zalcitabine

[Context]: Which among the following may be used as a sedatives-hypnotic?
A: Zolmitriptan
B: Zileuton
C: Zolpidem
D: Zalcitabine

Unperturbed

[Context]: Which among the following may be used as a sedatives-hypnotic?
A: Zolmitriptan
B: Zileuton
C: N-acetylglucosamine
D: Zalcitabine
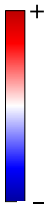
Adversarially perturbed