



MedSyn: LLM-Based Synthetic Medical Text Generation Framework

Gleb Kumichev¹(✉), Pavel Blinov², Yulia Kuzkina¹, Vasily Goncharov¹, Galina Zubkova², Nikolai Zenovkin¹, Aleksei Goncharov¹, and Andrey Savchenko²

¹ MIL Team, Moscow, Russia

{gleb.kumichev,yulia.kuzkina,vasily.goncharov,
nikolay.zenovkin}@mil-team.ru, alex.goncharov@mil-team.com

² Sber AI Lab, Moscow, Russia

{Blinov.P.D,GVZubkova,AVladSavchenko}@sber.ru

Abstract. Generating synthetic text addresses the challenge of data availability in privacy-sensitive domains such as healthcare. This study explores the applicability of synthetic data in real-world medical settings. We introduce MedSyn, a novel medical text generation framework that integrates large language models with a Medical Knowledge Graph (MKG). We use MKG to sample prior medical information for the prompt and generate synthetic clinical notes with GPT-4 and fine-tuned LLaMA models. We assess the benefit of synthetic data through application in the ICD code prediction task. Our research indicates that synthetic data can increase the classification accuracy of vital and challenging codes by up to 17.8% compared to settings without synthetic data. Furthermore, to provide new data for further research in the healthcare domain, we present the largest open-source synthetic dataset of clinical notes for the Russian language, comprising over 41k samples covering 219 ICD-10 codes.

Keywords: Synthetic data · Clinical note generation · ICD code prediction

1 Introduction

While extensive open medical datasets are available in English, like the MIMIC family of databases [18, 19] or the CPRD primary care database [14], their scope in comprehensively covering various medical areas is limited. The availability of textual medical data in non-English languages is even more constrained. Patient privacy and ethical considerations are major limiting factors to the public availability of such data. The latter remains a significant problem; the lack of textual medical resources substantially deters research, testing, and deployment of innovative Natural Language Processing (NLP) methods for national healthcare systems. Synthetic data generation addresses the issue of data scarcity in medical research.

Besides, the population’s diseases have a long-tail distribution, with rare diseases representing only a tiny fraction of cases in a dataset [24]. Such data imbalance problem directly affects the ML model’s performance on the downstream tasks [31, 38]. Since 2020, our clinical decision support system has been deployed in medical clinics in one of the regions. Insufficient text data on rare cases deters further system scaling, while synthetic (on-demand) medical note generation is the only solution.

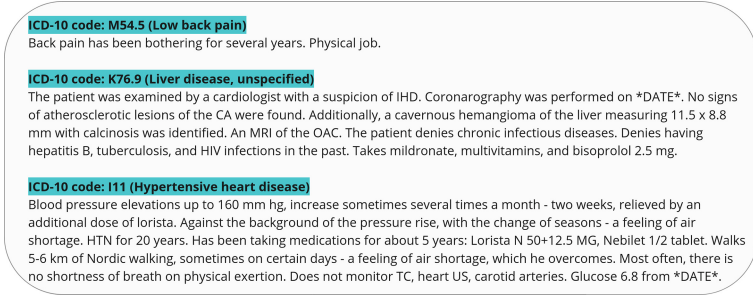


Fig. 1. Examples of real clinical notes from RuMedPrime dataset [35] (translated to English).

Nowadays, all patient information is stored in Electronic Health Records (EHRs), which contain a structured collection of medical events related to a patient and textual modality attributes: doctor’s clinical notes about symptoms and complaints, anamnesis, medication prescriptions, etc. Actual text from clinical notes is a complex object with typos, specialized terms, abbreviations, and contractions. Examples of such notes are shown in Fig. 1. That is why some early synthetic generation approaches (e.g. [10]) did not allow dealing with raw clinical text and tried to approximate EHRs only in terms of fixed categorical vectors and a limited set of factors, such as diagnosis and procedure codes or medication names. Including text fragments in synthetic EHRs has been challenging for a long time. Instead of generating medical text from scratch, some proposed frameworks heavily depend on real EHRs [8, 28], where a new health record is created by data imputation for some critical parts in the original one. However, such an approach limits the variability of results and leaves the risk of private data leakage.

The latest breakthroughs in developing Large Language Models (LLMs) open a new era in generating realistic, coherent, and diverse texts across various domains. Models like GPT-3 [7], LLaMA [37], and their successors have shown remarkable capabilities in text generation for general and medical texts [4]. However, even such powerful models still have some flaws [27]. First of all, they tend to make content errors and hallucinate [3], which is unacceptable in such a delicate area as medicine. Therefore, even LLM-based synthetic generation

frameworks still need external guidance and internal validation mechanisms to produce medically accurate and relevant texts.

Exploiting Medical Knowledge Graphs (MKGs) [12] and ontologies [1] is a way to mitigate the problem. Again, such resources are abundant for English but modestly available for less-represented languages like Russian. In this paper, we focus on developing a clinical note text generation framework combining LLMs capabilities with MKG in case studies for the Russian language.

Our key contributions can be summarized as follows:

1. We propose an open-source framework called MedSyn¹ for synthetic clinical note generation. The framework features a novel method that integrates disease-specific symptoms from an MKG and incorporates real data examples into the LLM generation pipeline to enhance the accuracy and diversity of generated data.
2. We introduce the first dataset² with synthetic clinical notes for the Russian language, which contains more than 41k clinical notes spanning over 219 ICD-10 (International Classification of Diseases) codes.
3. We provide results of experiments on synthetic data generation with the MedSyn framework, including comparisons between GPT-4 and open-sourced LLaMA-7b. It is shown that an open-sourced model fine-tuned on a specific dataset can perform on par with or surpass GPT-4's performance.

2 Related Work

2.1 Medical Knowledge Graphs

While a variety of MKGs exist in English [6, 9, 11, 40], few or none are available in other languages. There are different possibilities for MGK applications; for example, a line of work utilizes graph embeddings for various medical tasks like recommendation systems [13], NLI [33], and diagnosis prediction [43]. BioLORD [29] uses concepts and relationships from the knowledge graph as part of the LLM pre-training. Another approach for MKG utilization involves enriching the generation process with information extracted from these graphs. This strategy can be viewed as a specialized application of the retrieval-augmented generation framework [21], demonstrating the potential to produce more specific, diverse, and factually accurate language. However, applying such techniques in the medical domain is still an area that has not been extensively explored.

2.2 LLMs in Medical Domain

LLMs are increasingly utilized in the medical domain; they are primarily implemented for English [23, 34, 39] and Chinese [42, 45], evaluated for medical QA tasks, and used as medical chatbots. There is also a research direction that

¹ <https://github.com/milteam/MedSyn>.

² <https://huggingface.co/datasets/Glebkaa/MedSyn-synthetic>.

focuses on synthetic data generation. [26] trained the GPT-3 model from scratch using clinical and general English texts, then produced 20B of medical texts utilizing this model and introduced a smaller version of the model on synthetic data only. The resulting model outperforms ClinicalBERT [17] and the same model trained on actual data on MedNLI [30] and emrQA [25] benchmarks. The authors of [22] generated clinical texts and manually annotated them for the Named Entity Recognition (NER) task. The evaluation shows that the combination of original and synthetic corpora achieved better performance than using only the initial corpus. In [36], the authors improve performance on NER and relation extraction tasks with synthetic data, showing that increasing the number of synthetic sentences can improve model performance up to a certain point, beyond which the improvement becomes marginal. In a recent study [15], researchers explored the feasibility of using synthetic text as a training corpus for clinical NER in French. The findings suggest that synthetic clinical notes can be used to train NER models, although applications for other tasks remain to be explored.

The true potential of synthetic data in the medical field remains under active exploration [32, 36]. However, typical problems related to LLMs, like hallucinations, pose substantial challenges in such a critical field. Ensuring factual accuracy and addressing inconsistencies in medical models remain valuable concerns [41]. In our research, we strive to bridge the gap in controllable medical data generation, primarily focusing on the Russian language, which is heavily under-represented in linguistic medical resources.

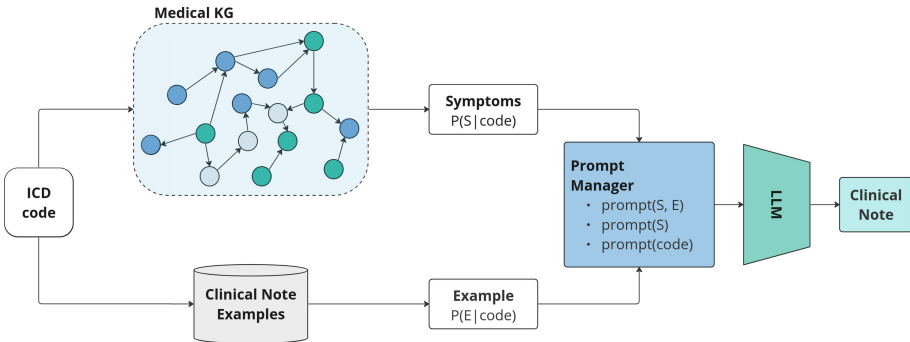


Fig. 2. The clinical notes generation pipeline implemented in MedSyn framework. Relevant symptoms from MKG and clinical note examples corresponding to the ICD code are compiled into a prompt and used as input for LLM inference.

3 Method

The overall pipeline for clinical note generation is illustrated in Fig. 2. To generate a clinical note for a target ICD code, data relevant to the MKG (Sect. 3.1)

and real examples are first sampled and combined into a prompt for LLM inference. We utilized GPT-4 and a fine-tuned LLaMA-7b for the LLMs (Russian Sect. 3.3). For fine-tuning LLaMA-7b, we constructed an instruction-following dataset (Russian Sect. 3.2). To generate a dataset of clinical notes for our experiments, we developed a specific generation task (Russian Sect. 3.4) with already prepared prompts.

3.1 Medical Knowledge Graph

As mentioned in Sect. 2.1, Russian-language equivalents of MKG are scarce. For our research, we used the WikiMed database as a foundation to develop the Russian MKG.

Table 1. MKG statistics. Di-Dr stands for disease-drug relation, Di-S for disease-symptom relation.

	Nodes			Edges	
	<i>Disease</i>	<i>Drug</i>	<i>Symptom</i>	<i>Di-Dr</i>	<i>Di-S</i>
#	2,747	2,968	2,554	1,997	2,554

The constructed MKG includes the following nodes: diseases (identified by ICD-10 codes), drugs, and symptoms. While diseases and drugs have predefined relations in this database, symptoms and their relations are not specified. The database includes clinical manifestations, which contain potential symptoms in a narrative format. To extract these symptoms, we utilized ChatGPT [2], prompting it to identify symptoms from the given text of clinical manifestations. For example, the clinical manifestation of tuberculosis, ‘*One of the common manifestations of spinal tuberculosis is the formation of cold abscesses on the neck and increased skin temperature*’, should lead to the extraction of symptoms [*cold abscesses on the neck, increased skin temperature*]. The extracted data were manually verified by comparing them with the initial text to ensure that only symptoms were included, and no irrelevant information or noise was extracted.

Finally extracted symptoms were then incorporated into the MKG. Its statistical details are presented in Table 1.

3.2 Instruction-Following Dataset

We collected a dataset of 152k Russian language samples focused on instruction-following for supervised fine-tuning³. These samples were derived from various medical benchmarks, databases, and the constructed MKG. Utilizing the MKG, we created questions that require multiple levels of reasoning, ranging from simple 1-hop to complex 3-hop distances. For example, a 1-hop reasoning question

³ <https://huggingface.co/datasets/Glebkaa/MedSyn-ift>.




Fig. 3. Examples of k-hop reasoning question on MKG. Di - Disease, Dr - Drug, S - Symptoms.

like ‘Provide symptoms for a disease’ directly connects diseases to symptoms (Di-S). A 2-hop question, such as ‘Write down medications that can be taken for these symptoms’, involves linking symptoms to diseases and then to drugs (S-Di-Dr). A more complex 3-hop reasoning question, like ‘List medications that can be taken for a disease if it is mistaken for another disease with similar symptoms’, maps diseases to symptoms, then to another disease, and finally to drugs (Di-S-Di-Dr), as shown in Fig. 3. We avoid more than three hops reasoning scenarios as, by our estimate, it produces too vague and error-prone samples. For the clinical notes, we employed two types of tasks: continuation, which extends an existing note from a random point, and generation, where a note is created from prior data like symptoms. We generated at least five different rephrasings for each to ensure instruction diversity.




Fig. 4. The structure of the instruction-following dataset. Leaves represent data sources and the percentage of data relative to the parent category.

In addition to real medical data, we also incorporate synthetic data from ChatGPT. Considering that real clinical notes often have many typos and stylistic variations, which may affect model performance, we suggest that adding synthetic notes could improve the model’s text generation and be a regularization method. To create this synthetic data, we prompted ChatGPT to generate clinical notes based on patient symptoms, age, and gender. For part of the data, style references with real samples were additionally provided. We also incorporate a medical dataset focused on typo correction to make the model more robust to typos. The structure of the dataset is represented in Fig. 4.

3.3 Fine-Tuning

Unlike the English language, to our knowledge, there aren't any open-source generative LLMs tailored for the medical domain in Russian. Thus, we employ GPT-4 [2] for data generation to establish a strong baseline.

Our work uses a model based on the LLaMA 2 family [37]. It is a collection of open generative language models with a parameter range from 7 to 70 billion. We fine-tuned the model with 7 billion parameters using a learning rate $2e^{-5}$ and a cosine learning rate scheduler to fine-tune the model. We utilized a global batch size of 256 and trained the model for three epochs.

To enhance the efficiency and accelerate the training of our model, we employed Low-Rank Adaptation (LoRA) [16]. This method involves freezing the model's weights and injecting trainable rank decomposition matrices into each layer of the Transformer architecture.

The pre-training data for LLaMa-7b consists of 90% English-language data and only 0.13% Russian-language data. Therefore, to fine-tune our model, we decided to use the pre-trained checkpoint from Saiga 2⁴ that is fine-tuned on Russian language instructions and dialogues generated by GPT-4.

3.4 Generation Task

We prepared a generation task to generate synthetic clinical notes with real data examples and symptoms spanning 105 ICD-10 category codes, as presented in the RuMedTop3 dataset [5]. We sample symptoms previously extracted from Russian MKG (Sect. 3.1) according to the approach outlined in Sect. 3.5.

We aim to achieve a uniform distribution of ICD codes for the generation task, but the lack of data requires inevitable trade-offs. Given the limited set of examples (1,283 samples), and to ensure that the sampling procedure represents the diversity of examples and symptoms, we have adopted a specific approach to determine the frequency of each ICD-10 category code \mathcal{C} and computes its weight based on the following rule:

$$w(\mathcal{C}) = J_3(N_{\text{symp}}^{\mathcal{C}}) \cdot J_3(N_{\text{exmp}}^{\mathcal{C}}), \quad (1)$$

where J_3 denotes the triple application of the function $J(x) = \log(1 + x)$, $N_{\text{exmp}}^{\mathcal{C}}$ refers to the number of examples corresponding to a given category code \mathcal{C} , and $N_{\text{symp}}^{\mathcal{C}}$ represents the number of all unique symptoms within that category. The logarithmic scale used in Eq. 1 is implemented to achieve a more uniform distribution of codes.

An exception to this weighting procedure is the category Z00, defined as *encounter for the general exam without complaint, suspect, or reported diagnosis*. As this category does not hold particular interest for downstream tasks, we set the number of generations for this category code to 10, thereby not factoring its weight into the overall distribution. We obtained the final generation task by sampling clinical notes and symptoms for this distribution, containing 2,503

⁴ https://huggingface.co/IlyaGusev/saiga2-7b_lora.

entries. Each entry consists of an ICD-10 code, an example of a real clinical note, and a subset of symptoms.

For the baseline, we generate samples that do not utilize data from MKG in their prompts. The baseline prompt is similar to the original one but contains only the disease name instead of incorporating disease prior information from MKG and a clinical note example. Generated and real clinical notes contain no ICD codes in the text to avoid data leaks.

3.5 Symptoms Sampling

The actual distribution of symptoms in clinical settings is complex. For example, certain symptoms may not coexist or be specific to a particular age or gender. In this study, however, we assume that symptoms are independently and identically distributed. Consequently, we select multiple symptoms for a disease without considering their inter-relationships. We randomly sample several symptoms from the MKG (Sect. 3.1) related to a disease, with the count ranging from 1 to 5, which is also chosen randomly.

3.6 Synthetic Dataset

We have released a dataset of 41,185 synthetic clinical notes in Russian, generated using GPT and fine-tuned LLaMA models spanning 219 ICD-10 codes. The dataset includes all generated samples, regardless of quality, to facilitate various data selection methods. More detailed statistics and descriptions of the data fields are provided in the project dataset repository.⁵ According to the provided licenses, all confidential information was anonymized, and researchers can safely use these datasets.

4 Experiments

4.1 Datasets and Tasks

In this research, we utilized the RuMedPrime dataset [35], containing 7,625 anonymized entries from outpatient visits to the Siberian State Medical University hospital. This dataset, unique as the only open-source collection of clinical notes in Russian annotated with ICD-10 codes, comprises each patient’s clinical note, symptoms, and corresponding ICD code. Based on this dataset the RuMedTop3 task was created, focusing on the ICD code prediction from a free-text clinical note. Given such a task, it is possible to implement an AI service that supports doctors with a second opinion on the diagnosis search.

Our study adopted the same dataset split as RuMedTop3, using 4,690 records for training, 848 for validation, and 822 for testing while incorporating full clinical notes alongside symptoms. Like RuMedTop3, we employ the second ICD-10 classification code hierarchy level. We also evaluated the results on the original RuMedTop3 dataset.

⁵ <https://huggingface.co/datasets/Glebkaa/MedSyn-synthetic>.

4.2 Models

We conducted experiments using both feature-based linear models and transformer models. For the linear model, we employed logistic regression based on term frequency-inverse document frequency (TF-IDF) features. For the transformer models, we run experiments with RuBERT [20] and RuBioRoBERTa [44] and report the average results from three runs.

4.3 Evaluation

ICD code prediction is a multi-class classification task. To evaluate it, we utilize the $hit@k$ score ($k \in [1, 3, 5]$), defined as follows:

$$hit@k = \frac{1}{N} \sum_{i=1}^N hit(\hat{y}, top_i^k), \quad (2)$$

where N is the number of samples and $hit(\hat{y}, top_i^k)$ is 1 if ground truth ICD code \hat{y} is on a ranked list of k predicted codes top_i^k and 0 otherwise.

4.4 Results

Prompt Following. We use the BERT-score [46] to measure the similarity of synthetic data to the examples and to the provided symptoms (Fig. 5).




Fig. 5. BERT-scores for example and symptoms usage.

As can be seen from the higher scores, the GPT-4 model follows instructions more precisely, produces results that are more similar to the example, and makes greater use of the provided symptoms.

While high similarity to the example is desirable, complete replication is unfavorable. To evaluate replication, we calculate the ratio of example N-grams usage, defined as the ratio of unique common N-grams between the generated sample and the example, divided by the number of unique N-grams in the example (Fig. 6). For most samples, the N-grams usage ratio is less than 1, suggesting that the examples are far from complete replication in the answer.




Fig. 6. Ratio of N-gram usage.




Fig. 7. The prediction results using only synthetic training data (codes K81 and I11). Contour bars represent the baseline prompt, which does not utilize MKG and consists solely of the task and the disease name.

Generating Data Out of the Training Set. One of the most exciting yet practically challenging scenarios involves generating data scarcely present in the original training set or generation of clinically valuable data. We selected two vital ICD codes for the experiment, K81 and I11. The first is cholecystitis, which affects about 20% of the adult population. The second code denotes a type of heart disease, one of the most common causes of death.

We transferred all real data samples to the test set, making evaluating the experiments with real data in the training set impossible. However, we prioritize a diverse test set in this experiment as it could mitigate the potential poor performance of unrepresented synthetic samples in downstream tasks. We replaced the real data in the training set with 30 synthetic samples for both models and added 59 samples for LLaMA-7b to assess the impact of scaling the number of samples (Fig. 7).

Although models trained with such synthetic data still have zero scores in the *hit@1* metric, they show promising results in less restricted metrics like *hit@5*, demonstrating the potential for further improvement in real data absence scenarios. Thus, synthetic data with specific refinements could increasingly become a viable alternative for training models in data-scarce environments.

Synthetic Upsampling. Another application for synthetic data is data upsampling. In this experiment, we used the same synthetic data as in the previous

section (Sect. 4.4) and added it to the training set. The results indicate that models can benefit from such synthetic data. For instance, the accuracy of K81 code prediction improves by 17.8% for the RuBioRoBERTa model (Fig. 8). To assess the overall accuracy across all ICD codes, we also evaluate both the baseline and the full prompts (Table 2).




Fig. 8. Results of prediction with upsampled training set for codes K81 and I11. The legend represents the data source/number of real samples/number of synthetic samples.

Table 2. Scores across all codes with upsampled training set for codes K81 and I11.

Model	Data	hit@1	hit@3	hit@5
Linear	Real	56.9	80.4	87.5
	LLaMA-7b-baseline	57.1	80.0	93.9
	LLaMA-7b	57.2	79.9	93.7
RuBioRoBERTa	Real	52.7	75.8	84.3
	LLaMA-7b-baseline	55.8	75.1	82.8
	LLaMA-7b	56.8	77.7	86.2

For a more detailed analysis, we focused on two codes that were frequently mistaken for each other more than any other pair. This decision was based on the confusion matrix, which measures how often each pair of codes is confused. The analysis revealed that the codes most often confused are M54 and G54.

We selected synthetic data for those codes generated via the same generation task described in Sect. 3.4 for the GPT-4 and LLaMA-7b models. For the LLaMA, we repeated the generation several times to evaluate the effect of data scaling. Here, we only report on the linear model to depict simultaneous changes for codes not averaged across several models. The experimental results are presented in Table 3. While data generated by GPT-4 provides improvements for both codes simultaneously, data generated by LLaMA still offers improvement for one of the codes without a drop for the other.

Table 3. Results of upsampling for the most pairwise misclassified codes. Prediction by the linear model. #R/S represents the number of real and synthetic samples in the training set. \uparrow represents growth of both codes simultaneously, \nearrow - growth for one code without drop for another.

Code	Data	# Real/Synthetic	hit@1	hit@3	hit@5
G54	Real	232/0	57.9	97.4	100
	GPT-4	232/14	60.5 \uparrow	97.4	100
	LLaMA-7b	232/14	57.9	97.4	100
	LLaMA-7b	232/72	60.5 \nearrow	97.4	100
M54	Real	560/0	85.1	98.9	100
	GPT-4	560/35	87.4 \uparrow	98.9	100
	LLaMA-7b	560/35	86.2 \nearrow	98.9	100
	LLaMA-7b	560/175	85.1	98.9	100

RuMedTop3 Upsampling. Although the generated clinical notes contain more information than the data in the RuMedTop3 task, which focuses on symptoms, using the generated data to upsample this dataset is still feasible, as they share the same set of ICD codes. We report results with generated data upsampling in Table 4, showing that all models benefit from the synthetic data.

Table 4. Results of upsampling on RuMedTop3 dataset (the real data size is 4,690 samples, and the size of the synthetic dataset is 2,503 samples).

Model	Data	hit@1	hit@3	hit@5
Linear	Real	49.8	72.7	87.8
	GPT-4	50.8	74.8	90.0
	LLaMA-7b	50.2	73.6	89.5
RuBERT	Real	46.5	70.4	79.3
	GPT-4	47.2	71.9	81.3
	LLaMA-7b	45.0	70.7	81.4
RuBioRoBERTa	Real	47.4	70.8	79.5
	GPT-4	47.3	71.7	80.4
	LLaMA-7b	46.1	70.4	79.6

4.5 Human Assessment

We performed the human evaluation in a side-by-side scenario to qualitatively assess the synthetic clinical texts. First, we randomly sampled 105 cases from real clinical notes examples according to the general ICD code distribution and paired

them with synthetic ones. Second, in each pair, we selected random sentences (with a median number of words of 8) to facilitate labeling and make a fair comparison detached from the notes structure. Such text pairs were presented to a medical intern with the only question – *Which text is generated, 1 or 2?* The assessor was correct in 58.09% (61 cases). Given that the random guessing is 50%, we can conclude that our synthetic texts have acceptable quality. In further research, we plan to evaluate the MedSyn framework in more elaborate human assessment scenarios.

5 Discussion

We used the generated datasets during all evaluations without applying filtration or sample selection techniques. Consequently, these datasets likely contain corrupted samples with minor factual errors or in some kind irrelevant to the provided prompt.

To estimate the validity of the samples, we predict their label using models trained on real data. We calculate the ratio of valid samples whose ground truth label appears in the top 5 predictions of at least 2 of five RuBERT models, each trained with different seeds. We found that 51% of LLaMA-7b samples and 64% of GPT-4 pass this criterion. However, this is only a coarse criterion as it may lead to false negatives, where a correct synthetic sample falls outside the training distribution and is consistently misclassified. Additionally, a sample might contain relevant information that leads to accurate predictions while still having some corruption. This observation also suggests that GPT-4 generated data might include fewer inaccurate samples, contributing to better performance. Possible sample corruption could lead to gaps in the authenticity and applicability of the generated content in specific clinical scenarios, highlighting the need for advanced filtration algorithms to refine the data quality. Future enhancements to the MKG, including a broader range of medical information, will likely improve the robustness and diversity of generated synthetic data.

While synthetic data is not directly tied to real patients, its use in clinical settings can still pose ethical questions regarding its applicability and acceptability. Key concerns include: 1) Ensuring that the data accurately reflects diverse patient populations without introducing biases; 2) Protecting against potential indirect privacy violations; 3) Assessing how its use might impact clinical decision-making. Additionally, it is essential to be transparent about how synthetic data is made and used and ensure its use follows informed consent rules in medical settings.

6 Conclusion

The proposed MedSyn framework suggests promising results in generating synthetic clinical notes. Human evaluation shows the high quality of generated texts, which are indistinguishable from real medical notes. In numerical experiments,

using additional synthetic notes leads to a 17.8% increase in ICD-code classification accuracy for vital and challenging classes compared to using only real data. Additionally, models trained on generated data reveal substantial quality even when used as the only training source, beating a solid baseline and helping to improve scores on the RuMedTop3 task. From a practical point of view, we plan to exploit the developed framework for rare disease note generation. Such synthetic data will allow us to substantially increase the number of disease classes in our clinical decision support system from tens to hundreds of ICD codes, giving the doctor a reliable second opinion even in rare scenarios.

The framework’s design allows easy integration with diverse MKGs, promising even more robust and varied data generation. To foster continued innovation in this field, we have made our trained model, part of the training dataset, and the synthetic dataset publicly available. These resources pave the way for further research in the medical field, especially in tasks where data is scarce. For instance, they potentially serve as datasets for medical NER tagging or in ICD coding tasks, where models trained on such data could provide valuable automated suggestions to humans. While synthetic data may contain inconsistencies or flaws, it is still precious in low-resource languages (like Russian) or low-data areas (like healthcare).

References

1. Abdollahi, M., Gao, X., Mei, Y., Ghosh, S., Li, J., Narag, M.: Substituting clinical features using synthetic medical phrases: Medical text data augmentation techniques. *Artifi. Intell. Med* **120**, 102167 (2021)
2. Achiam, J., et al.: GPT-4 technical report. arXiv preprint [arXiv:2303.08774](https://arxiv.org/abs/2303.08774) (2023)
3. Azaria, A., Mitchell, T.: The internal state of an llm knows when its lying. arXiv preprint [arXiv:2304.13734](https://arxiv.org/abs/2304.13734) (2023)
4. Benoit, J.R.: ChatGPT for clinical vignette generation, revision, and evaluation. *MedRxiv*, pp. 2023–02 (2023)
5. Blinov, P., Reshetnikova, A., Nesterov, A., Zubkova, G., Kokh, V.: RuMedBench: a Russian medical language understanding benchmark. In: *Artificial Intelligence in Medicine*, pp. 383–392. Springer International Publishing (2022). https://doi.org/10.1007/978-3-031-09342-5_38
6. Bodenreider, O.: The unified medical language system (umls): integrating biomedical terminology. *Nucleic Acids Res.* **32**(suppl_1), D267–D270 (2004)
7. Brown, T., et al.: Language models are few-shot learners. *Adv. Neural. Inf. Process. Syst.* **33**, 1877–1901 (2020)
8. Buczak, A.L., Babin, S., Moniz, L.: Data-driven approach for creating synthetic electronic medical records. *BMC Med. Inform. Decis. Mak.* **10**(1), 1–28 (2010)
9. Chandak, P., Huang, K., Zitnik, M.: Building a knowledge graph to enable precision medicine. *Sci. Data* **10**(1), 67 (2023)
10. Choi, E., Biswal, S., Malin, B., Duke, J., Stewart, W.F., Sun, J.: Generating multi-label discrete patient records using generative adversarial networks. In: *Proceedings of the 2nd Machine Learning for Healthcare Conference*, pp. 286–305. PMLR (2017)
11. Cui, H., et al.: A survey on knowledge graphs for healthcare: Resources, applications, and promises. arXiv preprint [arXiv:2306.04802](https://arxiv.org/abs/2306.04802) (2023)

12. Gao, Y., et al.: Leveraging a medical knowledge graph into large language models for diagnosis prediction. arXiv preprint [arXiv:2308.14321](https://arxiv.org/abs/2308.14321) (2023)
13. Gong, F., Wang, M., Wang, H., Wang, S., Liu, M.: SMR: medical knowledge graph embedding for safe medicine recommendation. *Big Data Res.* **23**, 100174 (2021)
14. Herrett, E., et al.: Data resource profile: clinical practice research datalink (CPRD). *Int. J. Epidemiol.* **44**(3), 827–836 (2015)
15. Hiebel, N., Ferret, O., Fort, K., Névéol, A.: Can synthetic text help clinical named entity recognition? a study of electronic health records in French. In: Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL), pp. 2320–2338. ACL, Dubrovnik, Croatia (May 2023)
16. Hu, E.J., et al.: Lora: Low-rank adaptation of large language models. arXiv preprint [arXiv:2106.09685](https://arxiv.org/abs/2106.09685) (2021)
17. Huang, K., Altosaar, J., Ranganath, R.: ClinicalBERT: Modeling clinical notes and predicting hospital readmission. arXiv preprint [arXiv:1904.05342](https://arxiv.org/abs/1904.05342) (2019)
18. Johnson, A.E., et al.: MIMIC-IV, a freely accessible electronic health record dataset. *Sci. Data* **10**(1), 1 (2023)
19. Johnson, A.E., et al.: MIMIC-III, a freely accessible critical care database. *Sci. Data* **3**(1), 1–9 (2016)
20. Kuratov, Y., Arkhipov, M.: Adaptation of deep bidirectional multilingual transformers for Russian language. arXiv preprint [arXiv:1905.07213](https://arxiv.org/abs/1905.07213) (2019)
21. Lewis, P., et al.: Retrieval-augmented generation for knowledge-intensive NLP tasks: Adv. Neural. Inf. Process. Syst. **33**, 9459–9474 (2020)
22. Li, J., et al.: Are synthetic clinical notes useful for real natural language processing tasks: A case study on clinical entity recognition. *J. Am. Med. Inform. Associat.* **28** (2021). <https://doi.org/10.1093/jamia/ocab112>
23. Luo, R., et al.: BioGPT: generative pre-trained transformer for biomedical text generation and mining. *Briefings Bioinform.* **23**(6) (2022)
24. Nguyen, T.T., et al.: Mimic-iv-icd: a new benchmark for extreme multilabel classification. arXiv preprint [arXiv:2304.13998](https://arxiv.org/abs/2304.13998) (2023)
25. Pampari, A., Raghavan, P., Liang, J., Peng, J.: emrQA: a large corpus for question answering on electronic medical records. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 2357–2368 (2018)
26. Peng, C., et al.: A study of generative large language model for medical research and healthcare. *NPJ Digital Med.* **6**(1), 210 (2023)
27. Ray, P.P.: ChatGPT: a comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope. *Internet of Things Cyber-Phys. Syst.* (2023)
28. Reiter, J.P., Drechsler, J.: Releasing multiply-imputed synthetic data generated in two stages to protect confidentiality. *Statistica Sinica*, 405–421 (2010)
29. Remy, F., Demuyne, K., Demeester, T.: BioLORD-2023: Semantic textual representations fusing llm and clinical knowledge graph insights. arXiv preprint [arXiv:2311.16075](https://arxiv.org/abs/2311.16075) (2023)
30. Romanov, A., Shivade, C.: Lessons from natural language inference in the clinical domain. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 1586–1596 (2018)
31. Santiso, S., Casillas, A., Pérez, A.: The class imbalance problem detecting adverse drug reactions in electronic health records. *Health Informatics J.* **25**(4), 1768–1778 (2019)
32. Shaib, C., Li, M.L., Joseph, S., Marshall, I.J., Li, J.J., Wallace, B.C.: Summarizing, simplifying, and synthesizing medical evidence using GPT-3 (with varying success). arXiv preprint [arXiv:2305.06299](https://arxiv.org/abs/2305.06299) (2023)

33. Sharma, S., Santra, B., Jana, A., Tokala, S., Ganguly, N., Goyal, P.: Incorporating domain knowledge into medical NLI using knowledge graphs. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). Association for Computational Linguistics (2019)
34. Singhal, K., et al.: Large language models encode clinical knowledge. *Nature* **620**(7972), 172–180 (2023)
35. Starovoytova, E., et al.: RuMedPrimeData (2021). <https://doi.org/10.5281/zenodo.5765873>
36. Tang, R., Han, X., Jiang, X., Hu, X.: Does synthetic data generation of LLMs help clinical text mining? arXiv preprint [arXiv:2303.04360](https://arxiv.org/abs/2303.04360) (2023)
37. Touvron, H., et al.: LLaMA: Open and efficient foundation language models. arXiv preprint [arXiv:2302.13971](https://arxiv.org/abs/2302.13971) (2023)
38. Wang, Y., Wei, Y., Yang, H., Li, J., Zhou, Y., Wu, Q.: Utilizing imbalanced electronic health records to predict acute kidney injury by ensemble learning and time series model. *BMC Med. Inform. Decis. Mak.* **20**, 1–13 (2020)
39. Wu, C., Zhang, X., Zhang, Y., Wang, Y., Xie, W.: PMC-LLaMA: Further finetuning llama on medical papers. arXiv preprint [arXiv:2304.14454](https://arxiv.org/abs/2304.14454) (2023)
40. Wu, X., Duan, J., Pan, Y., Li, M.: Medical knowledge graph: data sources, construction, reasoning, and applications. *Big Data Mining Analy.* **6**(2), 201–217 (2023). <https://doi.org/10.26599/BDMA.2022.9020021>
41. Xie, Q., Schenck, E.J., Yang, H.S., Chen, Y., Peng, Y., Wang, F.: Faithful AI in medicine: A systematic review with large language models and beyond. medRxiv (2023). <https://doi.org/10.1101/2023.04.18.23288752>
42. Xiong, H., et al.: DoctorGLM: Fine-tuning your chinese doctor is not a herculean task. arXiv preprint [arXiv:2304.01097](https://arxiv.org/abs/2304.01097) (2023)
43. Xu, X., et al.: Predictive modeling of clinical events with mutual enhancement between longitudinal patient records and medical knowledge graph. In: Proceedings of IEEE International Conference on Data Mining (ICDM), pp. 777–786 (2021)
44. Yalunin, A., Nesterov, A., Umerenkov, D.: RuBioRoBERTa: a pre-trained biomedical language model for Russian language biomedical text mining. arXiv preprint [arXiv:2204.03951](https://arxiv.org/abs/2204.03951) (2022)
45. Zhang, H., et al.: HuatuoGPT, towards taming language model to be a doctor. arXiv preprint [arXiv:2305.15075](https://arxiv.org/abs/2305.15075) (2023)
46. Zhang, T., Kishore, V., Wu, F., Weinberger, K.Q., Artzi, Y.: BERTScore: Evaluating text generation with BERT. arXiv preprint [arXiv:1904.09675](https://arxiv.org/abs/1904.09675) (2020)