

# High-Throughput Phenotyping of Clinical Text Using Large Language Models

Daniel B. Hier

*Dept. of Neurology & Rehabilitation  
University of Illinois at Chicago  
Chicago, IL USA  
dhier@uic.edu*

S. Ilyas Munzir

*Dept. of Neurology & Rehabilitation  
University of Illinois at Chicago  
Chicago, IL USA  
smunz2@uic.edu*

Anne Stahlfeld

*Dept. of Neurology & Rehabilitation  
University of Illinois at Chicago  
Chicago, IL USA  
astahl5@uic.edu*

Tayo Obafemi-Ajayi

*Engineering Program  
Missouri State University  
Springfield, MO USA  
tayoobafemiajayi@missouristate.edu*

Michael D. Carrithers

*Dept. of Neurology & Rehabilitation  
University of Illinois at Chicago  
Chicago, IL USA  
mcarl@uic.edu*

**Abstract**—High-throughput phenotyping automates the mapping of patient signs to standardized concepts, such as those in Human Phenotype Ontology (HPO), a process critical to precision medicine. We evaluated the automated phenotyping of clinical summaries from the Online Mendelian Inheritance in Man (OMIM) database using a large language model. Various APIs were used to automate text retrieval, sign identification, categorization, and normalization. GPT-4 outperformed GPT-3.5-Turbo in identifying, categorizing, and normalizing signs, achieving concordance with manual annotators comparable to concordance between manual annotators. While GPT-4 demonstrates high accuracy in sign identification and categorization, limitations remain in sign normalization, particularly in retrieving the correct HPO ID for a normalized term. Methods such as retrieval-augmented generation, changes in pre-training, and additional fine-tuning may help address these limitations. The combination of APIs with large language models presents a promising approach for high-throughput phenotyping of free text.

**Index Terms**—phenotype, large language model, natural language processing, high-throughput, OMIM, neurology, HPO, GPT-4.

## I. INTRODUCTION

Manual phenotyping of electronic health records is laborious and time-consuming [1], [2]. Precision medicine has driven the need for high-throughput phenotyping methods capable of processing large volumes of unstructured medical data efficiently [3], [4]. However, automating this process remains a challenge due to the complexity of medical text and the volume of physician notes [5]–[7]. Traditional natural language processing (NLP) methods for identifying phenotypic signs in clinical text have evolved from rule-based and dictionary-based systems [8], [9], to machine learning models [10], [11], and more recently to deep learning methods such as recurrent neural networks and convolutional neural networks [12]–[14]. Despite these advances, limitations remain, including low levels of accuracy, the need for large amounts of manually

annotated data to train models, and the inability to generalize models from one medical domain to another [2], [6], [7], [15].

The emergence of large language models (LLM) offers an opportunity to overcome some of these challenges, particularly for high-throughput phenotyping [16]–[18]. An LLM, such as GPT-4, is capable of understanding and generating human-like text across various domains due to pretraining on various data sources [19]. These models demonstrate strong zero-shot, one-shot, and few-shot learning abilities, allowing them to perform complex tasks such as extracting, categorizing, and normalizing clinical phenotypes without additional training [19]. Recent work [20], [21] has shown the potential of an LLM to automate the phenotyping process for large-scale electronic health records (EHR) and clinical summaries. An LLM can also derive phenotypes from other sources such as PubMed abstracts and clinical summaries [22].

Precision medicine relies on accurately computed patient phenotypes to guide treatment decisions and improve outcomes [23]. However, patient phenotypes recorded in EHRs are unstructured and require extraction, categorization, and normalization before they can be entered into precision medicine machine learning models. Human Phenotype Ontology (HPO) [1], [24], [25] is the most widely used standard to record phenotypic information, providing a structured vocabulary to describe the signs and symptoms of the disease. In the following, we refer to the signs and symptoms of the disease as *signs*. The Online Mendelian Inheritance in Man (OMIM) database organizes diseases into *phenotypic series*—a collection of diseases with similar clinical features but caused by mutations in different genes. For example, the dystonia phenotypic series includes DYT6, DYT11, and DYT25, all characterized by involuntary muscle contractions but with different underlying genetic causes [26], [27]. Automating the identification, categorization, and normalization of phenotypes from OMIM clinical summaries can serve as a useful surrogate for processing physician notes. This task benefits from the

TABLE I  
PHENOTYPIC SERIES PROCESSED BY HIGH-THROUGHPUT  
PHENOTYPING

Phenotypic Series	PS MIM	Diseases
amyotrophic lateral sclerosis (ALS)	PS105400	35
Charcot-Marie-Tooth disease (CMT)	PS118220	81
dystonia	PS128100	37
epilepsy generalized	PS600669	29
episodic ataxia	PS160120	9
familial febrile seizures	PS121210	17
hereditary spastic paraparesis (HSP)	PS303350	83
hyperekplexia	PS149400	4
leukodystrophy, hypomyelinating	PS312080	27
narcolepsy	PS161400	7
nemaline myopathy	PS161800	13
Parkinson	PS168600	33
progressive supranuclear palsy	PS601104	3
restless legs	PS102300	8
spinocerebellar ataxia	PS105400	40
striato nigral degeneration	PS609161	2

availability of OMIM text through an API and the absence of privacy regulations that govern patient data. Moreover, the phenotyping process in OMIM clinical summaries is similar to that required for physician notes [21].

In this study, we evaluated two large language models, specifically GPT-4 and GPT-3.5-Turbo, for high-throughput phenotyping of clinical text. By automating the process through APIs, we assess the capabilities of the models to identify, categorize, and normalize clinical signs. Furthermore, we visualize variability within a neurological disease phenotypic series using heat maps and dimension-reduced scatter plots, providing insights into the diversity of disease phenotypes.

## II. DATA

Neurological disease phenotypic data were retrieved from the OMIM database using the API (api.omim.org). For each disease within the OMIM database, the disease phenotypes are described in the *clinical synopsis* and the *clinical characteristics* sections. The *clinical synopsis* is a list of signs, symptoms, mode of inheritance, and age of onset, while the *clinical features* section summarizes published literature that underpins the phenotype of each disease. The OMIM API has separate calls for clinical features and clinical synopsis. Diseases in OMIM with similar phenotypes are grouped in a phenotypic series. OMIM currently has 582 phenotypic series, each with an identifier beginning with PS. We evaluated 16 phenotypic series that spanned across 405 neurogenetic diseases (Table I).

## III. METHODS

Figure 1 outlines the high-throughput phenotyping pipeline used for phenotype term extraction and normalization. Detailed parameters for the OMIM API and OpenAI API calls, along with the Python code and data files are available at the project’s GitHub site (<https://github.com/clslabMSU/highthroughput-phenotyping>).

**Text Extraction and Preprocessing.** Given a list of diseases and MIM numbers for each phenotypic series, the pipeline extracts clinical summaries from the OMIM API. White spaces

and tabs were converted to a single white space. Punctuation, including commas, hyphens, semicolons, single quotes, double quotes, forward slashes, and backslashes, were also standardized to a single white space. Periods were retained to identify sentence boundaries.

**Sign Identification.** The extracted text was passed to the OpenAI API with a structured prompt to identify neurological signs and symptoms (Box 1). This prompt format was designed to ensure consistency and clarity in extracting relevant signs while minimizing ambiguity.

**Sign Categorization.** The identified signs were categorized into 30 high-level categories (Box 2) using a subsequent OpenAI API call. The structured prompt was used to ensure that the signs were classified accurately into clinically relevant categories.

**Sign Normalization.** Signs were normalized by mapping them to the Human Phenotype Ontology (HPO) using two approaches. The first approach utilized spaCy (Explosion AI, Berlin) combined with Gensim BioWordVec embeddings. Vectors were generated for each sign and compared to HPO terms using cosine similarity, with the highest similarity assigned as the best match. The second approach involved GPT-4 or GPT-3.5-Turbo, where the models were tasked with mapping signs to HPO terms and IDs (Box 3).

### Box 1: Prompt for Sign Identification

You are a neurologist analyzing a case summary from OMIM. Your input is text containing ‘Clinical Features’ and ‘Description’. Extract relevant neurological symptoms (patient complaints) and signs (findings on examination). Here’s how the output should look:

‘Signs’: [‘symptom a’, ‘symptom b’, ‘symptom c’]

### Box 2: Prompt for Sign Categorization

You are a neurologist analyzing a list of signs. Classify each sign into one of these categories:

‘Behavior,’ ‘Bowel and Bladder,’ ‘Cognitive,’ ‘Deformity,’ ‘Dysautonomia,’ ‘Dystonia,’ ‘Extraocular Movements,’ ‘Fatigue,’ ‘Gait,’ ‘Head Shape,’ ‘Hearing,’ ‘Hyperkinesia,’ ‘Hyperreflexia,’ ‘Hypertonia,’ ‘Hypokinesia,’ ‘Hyporeflexia,’ ‘Hypotonia,’ ‘Incoordination,’ ‘Muscle Atrophy,’ ‘Other Cranial Nerve,’ ‘Pain,’ ‘Seizure,’ ‘Sensory,’ ‘Skin,’ ‘Sleep,’ ‘Speech,’ ‘Tremor,’ ‘Unclassified,’ ‘Vision,’ ‘Weakness.’

Your output should be a JSON object with each category as a key and a list of signs in that category as items.

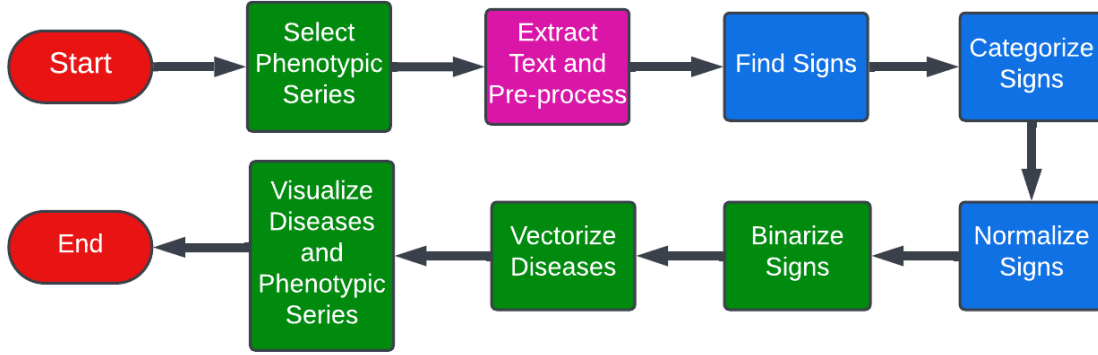


Fig. 1. **Pipeline for high-throughput phenotyping of clinical summaries from OMIM.** To support high-throughput, text retrieval, sign identification, sign categorization, and sign normalization is performed by an API.

### Box 3: Prompt for Sign Normalization

You are a neurologist tasked with mapping each sign to a concept in the Human Phenotype Ontology (HPO). Your output should be a JSON object with each input sign as a key and two item values: the ‘HPO Term’ and the ‘HPO ID.’

For example:

```
{‘input’: ‘Apraxia oral,’
‘HPO Term’: ‘Oromotor apraxia,’
‘HPO ID’: ‘HP:0000687’}
```

If the input term cannot be mapped to HPO, return ‘not-mappable’ in the ‘HPO Term’ and ‘HPO ID’ fields.

**Category Binarization.** To facilitate analysis, the 30 phenotype categories were binarized as either ‘0’ (no signs found in that category) or ‘1’ (one or more signs present). This binarization simplifies downstream analysis by reducing the data to presence/absence values, enabling easier comparisons of phenotypic similarities across diseases.

**Disease Vectorization.** For each disease, a vector was constructed from the 30 binary phenotype categories. Each element of the vector represents the presence (‘1’) or absence (‘0’) of a phenotype in the corresponding category. Among the 405 diseases evaluated, 283 had adequate clinical summaries for high-throughput phenotyping, and their disease vectors were stored as a data frame.

**Visualization of Disease Heterogeneity within a Phenotypic Series.** Heatmaps were created for each phenotypic series to visualize the heterogeneity of phenotypic presentations. Each row in the heatmap represents a disease, and each column represents one of the 30 binary phenotype categories (red for ‘present’, blue for ‘absent’).

**Visualization of Distances between the Centroids of Phenotypic Series.** Principal Component Analysis (PCA) was used

to reduce the 30 phenotype categories to two dimensions, allowing the visualization of distances between disease phenotypes in scatter plots (Fig. 7). Each centroid represents the phenotypic series’ average position, visualized as an ‘X’ on the scatter plot. The relative proximity between phenotypic series centroids highlights similarities between the diseases within a phenotypic series.

**Performance Metrics.** Disease processing rates, sign identification rates, sign categorization rates, and sign normalization rates were calculated based on 405 diseases, 175,724 words, and 16 phenotypic series (Table I). Sign identification, categorization, and normalization were validated using a dataset of 40 diseases from the Dystonia, Parkinson, Hereditary Spastic Paraparesis, and Charcot-Marie-Tooth phenotypic series.

TABLE II  
PERFORMANCE METRICS

Model	GPT-3.5 Turbo	GPT-4
Counts		
Diseases	405	405
Usable Diseases	207	283
Signs Identified	4,227	5,595
Unique Signs Identified	2,567	2,705
Words Processed	175,724	175,724
Rates†		
Disease Rate (sec/disease)	14.2	16.4
Identification Rate (sign/sec)	5.7	4.2
Categorization Rate (sign/sec)	2.9	2.3
Normalization Rate (sign/sec)	9.3	9.3

† Performance times and rates are representative. They were obtained on Apple Mac Studio with an M2 ultra CPU running Mac OS 14.5.

## IV. RESULTS

We performed high-throughput neurological phenotyping on 405 disease variants from 16 OMIM phenotypic series (Table I). Sign identification, sign categorization, and sign normalization were performed by GPT-3.5-Turbo or GPT-4 in three sequential submissions to the OpenAI API. The running




Fig. 2. **Heatmap for ALS phenotypic series with alphabetical category columns.** Diseases are along the y-axis. A unique MIM number identifies each disease. Compare to Fig. 3 with columns sorted by sign prevalence. Categories have been binarized so that 'red' indicates that the phenotype was present, and 'blue' indicates the phenotype was absent.




Fig. 3. **Heatmap for ALS phenotypic series with category columns sorted by sign prevalence** The most prevalent signs are weakness and muscle atrophy. Categories have been binarized so that 'red' indicates the phenotype was present and 'blue' indicates the phenotype was absent.

time per disease took 14.2s and 16.4s for GPT-3.5-Turbo and GPT-4 respectively. Although higher throughput might be possible with a faster CPU, more than 90% of the time expended was due to the four API calls.

The GPT-4 model outperformed the GPT-3.5-Turbo model on several performance metrics (Table II). GPT-4 produced usable data for 283 diseases, whereas GPT-3.5-Turbo produced usable data for 207 diseases. GPT-4 identified more signs (5,595 compared to 4,227) and more unique signs (2,705 compared to 2,567) than GPT-3.5-Turbo. The Jaccard Index, a stringent measure of concordance requiring exact matches between the large language models and the manual annotators, was higher for GPT-4 (0.31) than GPT-3.5-Turbo (0.16). A more relaxed measure of concordance, the maximum similarity index (based on cosine similarity from spaCy and BioWordVec embeddings from Gensim), showed high maximal mean similarities for signs compared to manual annotators (93.1 for GPT-3.5-Turbo and 94.2 for GPT-4). Weak matches (maximum similarity less than 0.80) were lower with GPT-4 than with GPT-3.5-Turbo. Compared to manual annotators, precision, recall, and F1 for sign identification were higher with GPT-4 than with GPT-3.5.




Fig. 4. **Heatmap for Charcot-Marie-Tooth phenotypic series.** The most prevalent signs are sensory symptoms, hyporeflexia, muscle atrophy, and weakness. MIM numbers for each disease in the phenotypic series are shown along the y-axis. Each row is a separate disease within the CMT phenotypic series and illustrates the diversity of phenotypic presentations of CMT within the phenotypic series. Categories have been binarized so that 'red' indicates the phenotype was present and 'blue' indicates the phenotype was absent.



Fig. 5. **Word cloud for phenotypic terms for Charcot-Marie-Tooth disease phenotypic series.** 939 Terms were identified through GPT-4 API. Term size reflects relative frequency. Note that many similar terms include 'areflexia', 'hyporeflexia', and 'decreased or absent reflexes'. Compare to Fig. 5 after terms have been further categorized by GPT-4 API.

The OpenAI API interface assigned each sign to one of 30 high-level categories. A significant simplification of the feature space was achieved by categorization of signs, as illustrated by comparing the word clouds for CMT signs (Fig.5 with CMT categories (Fig. 5). The ability of GPT-3.5-Turbo and GPT-4 to correctly assign signs to high-level categories was manually checked by a neurology expert for signs in the disease validation set. The accuracy of the GPT-4 was higher than that of the GPT-3.5-Turbo on sign categorization (94.0%



Fig. 6. **Word cloud for category frequencies for Charcot-Marie-Tooth (MCT) disease phenotypic series.** Phenotypic terms used to describe CMT diseases have been reduced to 30 categories. Word size in the word cloud reflects the size of each category. Compare to Fig. 4. The largest categories are Weakness, Deformity, and Gait.




Fig. 7. **Centroids plotted by phenotypic series.** The feature space has been reduced from 30 high-level categories to 2 dimensions by PCA. Each round marker is a disease in one of the five plotted phenotypic series. The **X** indicates the centroids for each phenotypic series. The expected proximities between ALS and HSP (both with weakness and spasticity) and between Parkinson and Dystonia (both movement disorders) are visualized. Five of the 16 available phenotypic series centroids are shown. Creating centroid plots for any combinations of the phenotypic series in Table I is possible. Due to concerns about interpretability, we have limited centroid plots to no more than 5 phenotypic series per plot.

compared to 58.4%). Sign categorization allowed us to create heat maps for each phenotypic series in which rows were diseases and columns were phenotype categories, as illustrated by Figs. 2 to Fig. 4. Distances between phenotypic series centroids can be plotted using PCA for dimension reduction. Fig. 7 shows an example of five centroids in a series of phenotypes. Sign normalization (Table V) was evaluated for the disease validation set. The SOTA NLP method performed best at 90.6% accuracy, followed by GPT-4 at 57.9% accuracy, and GPT-3.5-Turbo at 44.8% accuracy.

## V. DISCUSSION

We have developed a high-throughput pipeline that processes clinical text and identifies signs of disease. To support high-throughput, ease of use, and processing speed, the pipeline uses application programming interfaces (APIs) [28].

TABLE III  
SIGN IDENTIFICATION METRICS

Model	GPT-3.5-Turbo*	GPT-4*	Inter-Rater**
Signs Identified	358	609	694
Weak Matches (%)	15.0	11.6	4.0
Jaccard Index	0.16	0.35	0.36
Max Similarity Index	93.1	94.2	96.7
F1	0.52	0.66	0.60
Precision	0.61	0.66	0.96
Recall	0.45	0.65	0.44

\*Concordance for sign identification between GPT-3.5-Turbo and GPT-4 with the two manual annotators for 40 diseases in the validation dataset.

\*\*Inter-rater concordance for the manual annotators. Note that GPT-4 achieves a Jaccard Index similar to that between manual annotators.

TABLE IV  
SIGN CATEGORIZATION METRICS

Model	Accuracy	Precision	Recall
GPT-4	94.0	98.3	95.5
GPT-3.5 Turbo	58.4	78.4	95.2

Metrics based on manual review of sign categorization for the 40 diseases in the validation dataset.

We used an API to retrieve the summary text from OMIM and another API to allow GPT-4 to identify, categorize, and normalize signs. Clinical summaries from the OMIM database were utilized as our use case since the text is easily retrievable, rich with phenotypes, and not regulated as protected health information. However, these methods can be applied to text from other sources, including electronic health records, PubMed abstracts, full-text articles, and other clinical summaries.

Recognizing (identifying signs) and normalizing (mapping signs to an ontology) are challenging tasks for traditional NLP methods [2], [15], [29]–[32]. Progress has been made toward improving the recognition and normalization of medical concepts using transformers combined with specialized biomedical word embeddings [33], [34]. Large pre-trained language models provide a new approach to deep phenotyping (concept identification and normalization) that does not require additional training or a large corpus of manual annotations [22], [35]–[37]. Our pipeline for high-throughput phenotyping performed three phenotyping operations: sign identification, sign categorization, and sign normalization. In general, GPT-4 performed these operations with high accuracy and outperformed GPT-3.5-Turbo (Tables III, IV and V). Similarly, Groza et al. [22] evaluated GPT models for phenotype concept recognition using the ChatGPT interface. Their study demonstrated that GPT-4 outpaced the state-of-the-art methods in mention-level F1 scores of 0.7. Our work extends that of Groza et al. by demonstrating the utility of the GPT API to facilitate high-throughput phenotyping. In previous work, we have shown that GPT-4 can identify phenotypes in physician notes [20], [21], which is important for precision medicine [38], [39].

GPT-4 exhibited some weaknesses in sign normalization, achieving an accuracy of only 57.9%. This task has been noted by others as particularly challenging for GPT-4 [22]. In com-

TABLE V  
SIGN NORMALIZATION METRICS

Model	Accuracy	Precision	Recall
GPT-4	57.9	59.0	94.1
GPT-3.5 Turbo	44.8	49.8	52.9
SOTA NLP	90.6	90.8	99.8

Metrics based on manual review of the normalization of signs of the 40 diseases in the validation dataset. SOTA NLP is the spaCy cosine similarity method with Gensim BioWordVec embeddings.

parison, a state-of-the-art NLP model (SOTA) that combined BioWordVec from Gensim with the spaCy NLP similarity method demonstrated significantly higher accuracy at 90.6%. Although GPT-4 excelled at identifying plausible HPO terms for each input term, it was notably less accurate in providing the correct HPO IDs. In some instances, it even produced implausible HPO IDs. This discrepancy likely stems from GPT-4’s design, which relies heavily on pre-training to infer HPO IDs rather than employing a direct lookup capability. Currently, GPT-4 does not have an inherent mechanism to verify or retrieve accurate HPO IDs from a database. Shlyk et al. [40] have suggested remedying this limitation by adding retrieval augmented generation to assist in finding the correct HPO ID. Moreover, an inherent limitation of GPT models like GPT-4 is their non-deterministic nature. The choice of HPO ID for sign normalization can vary between different runs, even when the same input is provided [22]. This variability introduces inconsistencies that can be problematic in clinical applications where reliability is paramount.

We used GPT-4 to categorize the signs into 30 high-level categories. These high-level categories were chosen for their relevance to neurological phenotypes [41]. Although HPO has 28 high-level categories under *Phenotypic abnormality* [42], these categories are too broad to be useful in analyzing the phenotypes of neurological diseases. This categorization process significantly reduced the number of phenotypic terms needed to describe the diseases (compare Fig. 5 to Fig. 6). By assigning each phenotypic term to one of 30 high-level categories, we gained the ability to represent each disease in a phenotypic series as a row on a heatmap (Figs. 2 to 4). Heatmaps have also been used to visualize Orphadata disease phenotypes [41].

Once the phenotypic terms are acquired, a disease phenotype can be represented as a vector. Various methods are available to calculate the similarity between these disease vectors [43]–[48]. We used Principal Component Analysis (PCA) to reduce the dimensionality of these vectors to two dimensions (x and y), enabling us to visualize each disease as a marker on a scatter plot. To visualize the distances between the phenotypic series, we represent each series as a centroid of its component diseases. Although Fig. 7 is representative, these methods can be applied to display phenotypic distances between any combination of diseases or phenotypic series.

Large language models, including GPT-4, show promise for high-throughput phenotyping of clinical text, though some

issues identified in this work warrant further investigation. The level of accuracy required by an LLM for clinical decision-making remains uncertain [49]. It is important to recognize that human annotators do not always agree perfectly [50], and even expert physicians are susceptible to diagnostic errors [51]. There is a debate over whether health informatic tasks, such as phenotyping, are better suited to large general-purpose models or smaller, specially trained language models [52]. Concerns have been raised about the foundational weaknesses of large language models in healthcare, stemming from their limited training on EHR data [53]. Furthermore, an LLM struggles to process EHR data in tabular form (for example, the long tables of biochemical results) [54]. Groza et al. [22] have highlighted the stochastic nature of LLM outputs. If these models are to be used routinely in healthcare, issues of trust, privacy, equity, fairness, and confidentiality must be satisfactorily addressed [55], [56]. Furthermore, the problem of ‘hallucinations’ and ‘confabulations’ by LLMs remains unresolved [57]. The reduced accuracy of the LLM in retrieving the correct HPO ID (a normalization task) is a notable limitation (Table V).

This work has some limitations. While we tested GPT-3.5-Turbo and GPT-4, we did not compare their performance with other proprietary or open-source models. Future work will fully assess the robustness and error-handling capabilities of our pipeline. Scalability, cost analysis and stability studies are also required. Privacy concerns must still be addressed. More work is needed to better visualize disease phenotypes with heat maps and dimension-reduced plots. The generalizable of this pipeline to other disease domains such as cardiac, renal, hepatic, and rheumatic diseases should be explored. Limitations in retrieving the correct HPO ID needs to be addressed to assure that term normalization is performed with high accuracy.

Nonetheless, the case for applying LLMs to high-throughput phenotyping is compelling [18], [22], [35]–[37], [58]–[60]. These models are fast, accurate, and ready to run ‘out of the box.’ Unlike traditional neural network models, they do not rely on an extensive corpus of manual annotations. These models should be generalizable to a variety of diseases without additional training. Current limitations in sign normalization can be addressed using techniques from augmented retrieval generation [40], [61], by additional pre-training, or by creating small specialized models specifically for sign normalization. Large language models such as GPT-4 are expected to become the dominant method for high-throughput clinical text phenotyping.

## REFERENCES

- [1] P. N. Robinson, “Deep phenotyping for precision medicine,” *Human mutation*, vol. 33, no. 5, pp. 777–780, 2012.
- [2] J. Pathak, A. N. Kho, and J. C. Denny, “Electronic health records-driven phenotyping: challenges, recent advances, and perspectives,” *Journal of the American Medical Informatics Association*, vol. 20, no. e2, pp. e206–e211, 2013.
- [3] M. Afzal, S. R. Islam, M. Hussain, and S. Lee, “Precision medicine informatics: principles, prospects, and challenges,” *IEEE Access*, vol. 8, pp. 13 593–13 612, 2020.



- [4] M. Sahu, R. Gupta, R. K. Ambasta, and P. Kumar, "Artificial intelligence and machine learning in precision medicine: A paradigm shift in big data analysis," *Progress in Molecular Biology and Translational Science*, vol. 190, no. 1, pp. 57–100, 2022.
- [5] C. Shivade, P. Raghavan, E. Fosler-Lussier, P. J. Embi, N. Elhadad, S. B. Johnson, and A. M. Lai, "A review of approaches to identifying patient phenotype cohorts using electronic health records," *Journal of the American Medical Informatics Association*, vol. 21, no. 2, pp. 221–230, 2014.
- [6] Y. Chen, R. J. Carroll, E. R. M. Hinz, A. Shah, A. E. Eyler, J. C. Denny, and H. Xu, "Applying active learning to high-throughput phenotyping algorithms for electronic health records data," *Journal of the American Medical Informatics Association*, vol. 20, no. e2, pp. e253–e259, 2013.
- [7] K. P. Liao, J. Sun, T. A. Cai, N. Link, C. Hong, J. Huang, J. E. Huffman, J. Gronsbell, Y. Zhang, Y.-L. Ho *et al.*, "High-throughput multimodal automated phenotyping (map) with application to phewas," *Journal of the American Medical Informatics Association*, vol. 26, no. 11, pp. 1255–1262, 2019.
- [8] S. Eltyeb and N. Salim, "Chemical named entities recognition: a review on approaches and applications," *Journal of cheminformatics*, vol. 6, no. 1, pp. 1–12, 2014.
- [9] A. P. Quimbaya, A. S. Múnera, R. A. G. Rivera, J. C. D. Rodríguez, O. M. M. Velandia, A. A. G. Peña, and C. Labbé, "Named entity recognition over electronic health records through a combined dictionary-based approach," *Procedia Computer Science*, vol. 100, pp. 55–61, 2016.
- [10] L. Hirschman, A. A. Morgan, and A. S. Yeh, "Rutabaga by any other name: extracting biological names," *Journal of Biomedical Informatics*, vol. 35, no. 4, pp. 247–259, 2002.
- [11] Ö. Uzuner, B. R. South, S. Shen, and S. L. DuVall, "2010 i2b2/va challenge on concepts, assertions, and relations in clinical text," *Journal of the American Medical Informatics Association*, vol. 18, no. 5, pp. 552–556, 2011.
- [12] G. Lample, M. Ballesteros, S. Subramanian, K. Kawakami, and C. Dyer, "Neural architectures for named entity recognition," *arXiv preprint arXiv:1603.01360*, 2016.
- [13] J. P. Chiu and E. Nichols, "Named entity recognition with bidirectional LSTM-CNNs," *Transactions of the Association for Computational Linguistics*, vol. 4, pp. 357–370, 2016.
- [14] M. Habibi, L. Weber, M. Neves, D. L. Wiegandt, and U. Leser, "Deep learning with word embeddings improves biomedical named entity recognition," *Bioinformatics*, vol. 33, no. 14, pp. i37–i48, 2017.
- [15] H. Alzoubi, R. Alzubi, N. Ramzan, D. West, T. Al-Hadhrani, and M. Alazab, "A review of automatic phenotyping approaches using electronic health records," *Electronics*, vol. 8, no. 11, p. 1235, 2019.
- [16] C. Yan, H. Ong, M. Grabowska, M. Krantz, W.-C. Su, A. Dickson, J. F. Peterson, Q. Feng, D. M. Roden, C. M. Stein *et al.*, "Large language models facilitate the generation of electronic health record phenotyping algorithms," *medRxiv*, pp. 2023–12, 2023.
- [17] J. Yang, C. Liu, W. Deng, D. Wu, C. Weng, Y. Zhou, and K. Wang, "Enhancing phenotype recognition in clinical notes using large language models: Phenobcber and phenogpt," *Patterns*, 2023.
- [18] A. Wang, C. Liu, J. Yang, and C. Weng, "Fine-tuning large language models for rare disease concept normalization," *bioRxiv*, pp. 2023–12, 2023.
- [19] J. Qiu, L. Li, J. Sun, J. Peng, P. Shi, R. Zhang, Y. Dong, K. Lam, F. P.-W. Lo, B. Xiao *et al.*, "Large ai models in health informatics: Applications, challenges, and the future," *IEEE Journal of Biomedical and Health Informatics*, 2023.
- [20] S. I. Munzir, D. B. Hier, C. Oommen, and M. D. Carrithers, "A large language model outperforms other computational approaches to the high-throughput phenotyping of physician notes," *arXiv preprint arXiv:2406.14757*, 2024, accepted in AMIA Annual Symposium 2024. [Online]. Available: <https://arxiv.org/abs/2406.14757>
- [21] S. I. Munzir, D. B. Hier, and M. D. Carrithers, "High throughput phenotyping of physician notes with large language and hybrid nlp models," *arXiv preprint arXiv:2403.05920*, 2024, accepted in International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC 2024). [Online]. Available: <https://doi.org/10.48550/arXiv.2403.05920>
- [22] T. Groza, H. Caufield, D. Gratton, G. Baynam, M. A. Haendel, P. N. Robinson, C. J. Mungall, and J. T. Reese, "An evaluation of gpt models for phenotype concept recognition," *BMC Medical Informatics and Decision Making*, vol. 24, no. 1, p. 30, 2024.
- [23] J. P. Ackerman, D. C. Bartos, J. D. Kapplinger, D. J. Tester, B. P. Delisle, and M. J. Ackerman, "The promise and peril of precision medicine: phenotyping still matters most," in *Mayo Clinic Proceedings*, vol. 91, no. 11. Elsevier, 2016, pp. 1606–1616.
- [24] J. S. Amberger, C. A. Bocchini, F. Schiettecatte, A. F. Scott, and A. Hamosh, "Omim. org: Online mendelian inheritance in man (omim®), an online catalog of human genes and genetic disorders," *Nucleic acids research*, vol. 43, no. D1, pp. D789–D798, 2015.
- [25] S. Köhler, N. A. Vasilevsky, M. Engelstad, E. Foster, J. McMurry, S. Aymé, G. Baynam, S. M. Bello, C. F. Boerkoel, K. M. Boycott *et al.*, "The human phenotype ontology in 2017," *Nucleic acids research*, vol. 45, no. D1, pp. D865–D876, 2017.
- [26] D. Hier, R. Yelugam, S. Azizi, and D. Wunsch III, "A focused review of deep phenotyping with examples from neurology," *Eur Sci J*, vol. 18, pp. 4–19, 2022.
- [27] D. Hier, R. Yelugam, S. Azizi, M. Carrithers, and I. Wunsch, "Dc. high throughput neurological phenotyping with metamap," *Eur Sci J*, vol. 18, pp. 37–49, 2022.
- [28] A. Tarkowska, D. Carvalho-Silva, C. E. Cook, E. Turner, R. D. Finn, and A. D. Yates, "Eleven quick tips to build a usable rest api for life sciences," *PLoS computational biology*, vol. 14, no. 12, p. e1006542, 2018.
- [29] S. Fu, D. Chen, H. He, S. Liu, S. Moon, K. J. Peterson, F. Shen, L. Wang, Y. Wang, A. Wen *et al.*, "Clinical concept extraction: a methodology review," *Journal of Biomedical Informatics*, p. 103526, 2020.
- [30] M. Agrawal, C. O'Connell, Y. Fatemi, A. Levy, and D. Sontag, "Robust benchmarking for machine learning of clinical entity extraction," in *Machine Learning for Healthcare Conference*. PMLR, 2020, pp. 928–949.
- [31] Y. Zhang, T. Cai, S. Yu, K. Cho, C. Hong, J. Sun, J. Huang, Y.-L. Ho, A. N. Ananthakrishnan, Z. Xia *et al.*, "High-throughput phenotyping with electronic medical record data using a common semi-supervised approach (phecap)," *Nature protocols*, vol. 14, no. 12, pp. 3426–3444, 2019.
- [32] S. Yang, P. Varghese, E. Stephenson, K. Tu, and J. Gronsbell, "Machine learning approaches for electronic health records phenotyping: a methodical review," *Journal of the American Medical Informatics Association*, vol. 30, no. 2, pp. 367–381, 2023.
- [33] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang, "BioBERT: a pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2020.
- [34] Z. Ji, Q. Wei, and H. Xu, "Bert-based ranking for biomedical entity normalization," *AMIA Summits on Translational Science Proceedings*, vol. 2020, p. 269, 2020.
- [35] C. Shyr, Y. Hu, P. A. Harris, and H. Xu, "Identifying and extracting rare disease phenotypes with large language models," *arXiv preprint arXiv:2306.12656*, 2023.
- [36] C. Shyr, Y. Hu, L. Bastarache, A. Cheng, R. Hamid, P. Harris, and H. Xu, "Identifying and extracting rare diseases and their phenotypes with large language models," *Journal of Healthcare Informatics Research*, pp. 1–24, 2024.
- [37] N. J. Dobbins, "Generalizable and scalable multistage biomedical concept normalization leveraging large language models," *arXiv preprint arXiv:2405.15122*, 2024.
- [38] M. Simmons, A. Singhal, and Z. Lu, "Text mining for precision medicine: bringing structure to ehrs and biomedical literature to understand genes and health," *Translational Biomedical Informatics: A Precision Medicine Perspective*, pp. 139–166, 2016.
- [39] A. Sitapati, H. Kim, B. Berkovich, R. Marmor, S. Singh, R. El-Kareh, B. Clay, and L. Ohno-Machado, "Integrated precision medicine: the role of electronic health records in delivering personalized treatment," *Wiley Interdisciplinary Reviews: Systems Biology and Medicine*, vol. 9, no. 3, p. e1378, 2017.
- [40] D. Shlyk, T. Groza, M. Mesiti, S. Montanelli, and E. Cavalleri, "Real: A retrieval-augmented entity linking approach for biomedical concept recognition," in *Proceedings of the 23rd Workshop on Biomedical Natural Language Processing*, 2024, pp. 380–389.
- [41] D. B. Hier, R. Yelugam, M. D. Carrithers, and D. C. Wunsch III, "The visualization of orphadata neurology phenotypes," *Frontiers in Digital Health*, vol. 5, p. 1064936, 2023.
- [42] J. Foreman, S. Brent, D. Perrett, A. P. Bevan, S. E. Hunt, F. Cunningham, M. E. Hurlles, and H. V. Firth, "Decipher: Supporting the interpretation and sharing of rare disease phenotype-linked variant data to advance

diagnosis and research,” *Human Mutation*, vol. 43, no. 6, pp. 682–697, 2022.

- [43] T. Mabotuwana, M. C. Lee, and E. V. Cohen-Solal, “An ontology-based similarity measure for biomedical data—application to radiology reports,” *Journal of biomedical informatics*, vol. 46, no. 5, pp. 857–868, 2013.
- [44] L. Cheng, H. Zhao, P. Wang, W. Zhou, M. Luo, T. Li, J. Han, S. Liu, and Q. Jiang, “Computational methods for identifying similar diseases,” *Molecular Therapy-Nucleic Acids*, vol. 18, pp. 590–604, 2019.
- [45] A. Gamba, M. Salmona, L. Cantù, and G. Bazzoni, “The similarity of inherited diseases (ii): clinical and biological similarity between the phenotypic series,” *BMC Medical Genomics*, vol. 13, pp. 1–11, 2020.
- [46] A. Gamba, M. Salmona, and G. Bazzoni, “The similarity of inherited diseases (i): clinical similarity within the phenotypic series,” *BMC Medical Genomics*, vol. 14, pp. 1–12, 2021.
- [47] M. Chagoyen and F. Pazos, “Characterization of clinical signs in the human interactome,” *Bioinformatics*, vol. 32, no. 12, pp. 1761–1765, 2016.
- [48] H. Xue, J. Peng, and X. Shang, “Predicting disease-related phenotypes using an integrated phenotype similarity measurement based on hpo,” *BMC systems biology*, vol. 13, pp. 1–12, 2019.
- [49] Z. Grotenhuis, “Text mining of clinical outcomes for medical research: how accurate should it be?” MSc. Thesis, Utrecht University, Utrecht, The Netherlands, October 2022.
- [50] C. Oommen, Q. Howlett-Prieto, M. D. Carrithers, and D. B. Hier, “Inter-rater agreement for the annotation of neurologic signs and symptoms in electronic health records,” *Frontiers in Digital Health*, vol. 5, p. 1075771, 2023.
- [51] M. I. Chimowitz, E. L. Logigian, and L. R. Caplan, “The accuracy of bedside neurological diagnoses,” *Annals of neurology*, vol. 28, no. 1, pp. 78–85, 1990.
- [52] E. Hernandez, D. Mahajan, J. Wulff, M. J. Smith, Z. Ziegler, D. Nadler, P. Szolovits, A. Johnson, E. Alsentzer *et al.*, “Do we still need clinical language models?” in *Conference on Health, Inference, and Learning*. PMLR, 2023, pp. 578–597.
- [53] M. Wornow, Y. Xu, R. Thapa, B. Patel, E. Steinberg, S. Fleming, M. A. Pfeffer, J. Fries, and N. H. Shah, “The shaky foundations of large language models and foundation models for electronic health records,” *npj Digital Medicine*, vol. 6, no. 1, p. 135, 2023.
- [54] J. Lovón-Melgarejo, T. Ben-Haddi, J. Di Scala, J. G. Moreno, and L. Tamine, “Revisiting the mimic-iv benchmark: Experiments using language models for electronic health records,” in *Proceedings of the First Workshop on Patient-Oriented Language Processing (CL4Health)@ LREC-COLING 2024*, 2024, pp. 189–196.
- [55] S. Harrer, “Attention is not all you need: the complicated case of ethically using large language models in healthcare and medicine,” *EBioMedicine*, vol. 90, 2023.
- [56] L. Sun, Y. Huang, H. Wang, S. Wu, Q. Zhang, C. Gao, Y. Huang, W. Lyu, Y. Zhang, X. Li *et al.*, “Trustllm: Trustworthiness in large language models,” *arXiv preprint arXiv:2401.05561*, 2024.
- [57] I. S. Schwartz, K. E. Link, R. Daneshjou, and N. Cortés-Penfield, “Black box warning: large language models and the future of infectious diseases consultation,” *Clinical infectious diseases*, vol. 78, no. 4, pp. 860–866, 2024.
- [58] W. E. Thompson, D. M. Vidmar, J. K. De Freitas, J. M. Pfeifer, B. K. Fornwalt, R. Chen, G. Altay, K. Manghnani, A. C. Nelsen, K. Morland *et al.*, “Large language models with retrieval-augmented generation for zero-shot disease phenotyping,” *arXiv preprint arXiv:2312.06457*, 2023.
- [59] Ş. Kafkas, M. Abdelhakim, A. Althagafi, S. Toonsi, M. Alghamdi, P. N. Schofield, and R. Hoehndorf, “The application of large language models to the phenotype-based prioritization of causative genes in rare disease patients,” *medRxiv*, pp. 2023–11, 2023.
- [60] A. Wang, C. Liu, J. Yang, and C. Weng, “Fine-tuning large language models for rare disease concept normalization,” *Journal of the American Medical Informatics Association*, p. ocae133, 2024.
- [61] J. Chen, H. Lin, X. Han, and L. Sun, “Benchmarking large language models in retrieval-augmented generation,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 17 754–17 762.