

ReToP: Apprendre à réécrire les dossiers médicaux de patients pour la prédiction clinique

Jesús Lovón-Melgarejo¹ Jose G. Moreno¹
Christine Damase-Michel² Lynda Tamine¹

(1) Université de Toulouse, IRIT UMR 5505, Toulouse, France

(2) Centre Hospitalier Universitaire de Toulouse,
CERPOP INSERM UMR 1295 - équipe SPHERE,

Faculté de Médecine Université de Toulouse, Toulouse, France

(1) first.last@irit.fr, (2) first.last@utoulouse.fr

RÉSUMÉ

Bien que les dossiers médicaux de patients (DMP) fournissent des informations cliniques essentielles, leur complexité rend la modélisation prédictive difficile (Lovón-Melgarejo *et al.*, 2025). Les grands modèles de langue (GML) atténuent en partie ces limitations en enrichissant les représentations des DMP grâce à des connaissances médicales. Toutefois, les méthodes sous-jacentes sont indépendantes des tâches et n'exploitent pas pleinement les signaux spécifiques aux objectifs prédictifs de ces tâches. Dans cet article, nous présentons **Rewrite-To-Predict (ReToP)**, une approche fondée sur des LLM qui réalise un entraînement bout-à-bout d'un modèle de réécriture de DMP et d'un modèle de prédiction clinique. **ReToP** réalise un alignement de la réécriture avec la prédiction via un score *Classifier Supervised Contribution (CSC)* permettant la génération de réécritures cliniquement pertinentes qui renforce directement la performance prédictive. Sur trois tâches cliniques issues de MIMIC-IV, **ReToP** dépasse les modèles de référence et démontre une forte généralisabilité à de nouveaux ensembles de données et tâches avec un ajustement minimal. *Ce travail est paru dans les actes de la conférence Web Search and Data Mining Conference WSDM '26 et sera présenté à l'atelier LLM@Hopital'26. Pour une description détaillée se référer à l'article Lovón-Melgarejo *et al.* (2026).*

ABSTRACT

ReToP : Learning to Rewrite Electronic Health Records for Clinical Prediction

Although Electronic Health Records (EHRs) provide crucial clinical information, their complexity challenges predictive modeling. Large Language Models (LLMs) mitigate these issues by enhancing EHR representation with medical knowledge. However, existing approaches are mainly task-agnostic and do not adequately integrate task-specific signals. In this work, we propose **Rewrite-To-Predict (ReToP)**, an LLM-based framework that addresses this limitation through an end-to-end training of an EHR rewriter and a clinical predictor. **ReToP** aligns the rewriter with prediction objectives using a novel *Classifier Supervised Contribution (CSC)* score that enables the rewriter to generate clinically relevant rewrites that enhance prediction. Our **ReToP** framework surpasses strong baseline across three clinical tasks on MIMIC-IV and demonstrates generalizability to unseen datasets and tasks with minimal fine-tuning while preserving faithful rewrites and emphasizing task-relevant features. *This work has been published in proceedings of Web Search and Data Mining Conference WSDM '26 and will be presented at the workshop LLM@Hopital'26. For a detailed description of the work, see Lovón-Melgarejo *et al.* (2026)*

1 Contexte et défis scientifiques

Les modèles prédictifs cliniques s'appuient fortement sur les dossiers médicaux électroniques de Patients (DMP), qui codifient les caractéristiques médicales longitudinales des patients (par exemple, maladies, procédures) afin d'estimer les risques de développer ou de présenter une pathologie. La prédiction clinique est confrontée à plusieurs défis, parmi lesquels l'hétérogénéité, la haute dimensionnalité et l'aspect "creux" des DME (Landi *et al.*, 2020; Zhu *et al.*, 2024). D'importants efforts de recherche ont été consacrés à relever ces défis grâce au développement de modèles d'apprentissage automatique pour un large éventail de tâches de prédiction clinique, notamment la prédiction du diagnostic (Madhumita *et al.*, 2018; Dligach & Miller, 2018; Hur *et al.*, 2023; Kim *et al.*, 2024), la prédiction de la mortalité (Choi *et al.*, 2020; Madhumita *et al.*, 2018; Hur *et al.*, 2023; Naik *et al.*, 2022), la prédiction des réadmissions (Jiang *et al.*, 2024; Hur *et al.*, 2023) et la prédiction de la durée du séjour (K *et al.*, 2022; Hur *et al.*, 2023). L'une des principales conclusions de ces études est que l'exploitation des connaissances d'experts associées aux informations issues des données des patients améliore la modélisation des DSE, ce qui se traduit par une amélioration significative de la précision des prédictions. Les études précédentes intégraient principalement les connaissances symboliques des experts sous la forme de graphes de connaissances cliniques représentant les caractéristiques médicales (par exemple, les codes de diagnostic CIM-10) et leurs relations afin de modéliser des informations contextuelles utiles pour compléter les données des DSE (Choi *et al.*, 2018, 2020; Tamine & Goeuriot, 2021; Xu *et al.*, 2024).

Par la suite, les grands modèles de langue (LLM) sont apparus comme des modèles marquant la transition des connaissances symboliques vers les connaissances paramétriques, nous rapprochant ainsi des graphes de connaissances potentiels dotés de capacités impressionnantes de compréhension et de génération du langage (Petroni *et al.*, 2019). À ce jour, de nombreux travaux démontrent le potentiel des LLMs pour effectuer des tâches médicales (Wu *et al.*, 2024; Li *et al.*, 2024; Kweon *et al.*, 2024). Dans la littérature relative à la prédiction clinique basée sur les LLMs, les modèles s'appuient sur l'une des approches suivantes illustrées dans la figure 1 : (a) l'utilisation de LLMs pré-entraînés et affinés, spécialisés dans les corpus textuels médicaux (par exemple, Llamacare (Li *et al.*, 2024), Llemr (Wu *et al.*, 2024)). Bien qu'ils démontrent de solides capacités d'interprétation des caractéristiques complexes des DMP, ces travaux ont été évalués en particulier sur des tâches de compréhension des connaissances médicales (par exemple, la réponse à des questions médicales) et nécessitent généralement un coût d'ajustement important pour atteindre une capacité de généralisation (Liu *et al.*, 2024b) ; (b) compléter les DMP avec des connaissances cliniques issues de LLMs, soit de manière agnostique (Xu *et al.*, 2024), soit en fonction des données des DMP avant d'entraîner le modèle prédictif (Jiang *et al.*, 2024).

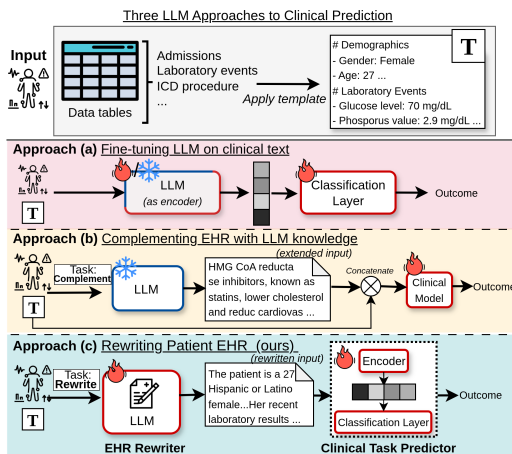


FIGURE 1 – Comparaison de trois approches basées LLMs pour la prédiction clinique à partir des DMP. Approche (a) fondées sur des LLMs comme des encoders de DMP. Approche (b) fondée sur des LLMs qui complètent des DMP. Notre approche, ReToP (c), utilise un LLM entraîné pour la réécriture des DMP de façon alignée alignée avec la tâche extrinsèque de prédiction clinique.

Cependant, ces travaux existants présentent deux limites principales :

1. Ils offrent une faible flexibilité pour déterminer dans quelle mesure une connaissance est utile ou bruitée lors de la fusion naïve entre les connaissances paramétriques des LLMs et les connaissances KG (Chen *et al.*, 2024; Liu *et al.*, 2024a) ;
2. Ils ne relient pas la complétion des DMP à la prédiction clinique. Étant donné que les graphes de connaissances contiennent une grande quantité d'informations relationnelles, compléter les DMP avec des données bruitées peut nuire à la précision des prédictions.

Inspirés par la manière dont les experts humains établissent leurs diagnostics, nous nous posons une question clé : « *Compte tenu du DMP du patient, quels sont les fondements cliniques qui justifient la prédiction ?* » Notre travail répond à cette question tout en apportant des solutions aux limites susmentionnées.

Nous proposons un modèle de réécriture de DMP basé sur le LLM, qui agit comme source unique de connaissances expertes, capable de générer des descriptions fidèles et pertinentes du DMP original du patient. Le module de réécriture du DMP basé sur le LLM est ensuite affiné afin de réviser les réécritures du DMP pour s'aligner avec la prédiction clinique. Cependant, la mise en œuvre de cette réponse pose les deux défis suivants :

- **C1** : À partir d'un DMP patient, le modèle de réécriture basé sur un LLM pourrait être capable de rappeler fidèlement, parmi un grand nombre de caractéristiques cliniques, les plus pertinentes pour toute tâche de prédiction. Une solution potentielle consiste à utiliser un module de réécriture de DMP basé sur un LLM disponible. Cependant, des travaux antérieurs ont montré que les LLM ont encore du mal à comprendre les données tabulaires, y compris

les DMP (Sui *et al.*, 2024; Singha *et al.*, 2023; Lovón-Melgarejo *et al.*, 2025). Étant donné que les caractéristiques cliniques pertinentes d'un DMP sont a priori inconnues, l'alternative consistant à concevoir un LLM entraînable pour la réécriture de DMP pose le défi d'un manque de supervision de référence.

- **C2** : Les caractéristiques cliniques pertinentes des DMP sont susceptibles de dépendre de la tâche. Ainsi, la précision des prédictions pourrait servir d'évaluation indirecte de la qualité de la réécriture des DMP. Cependant, la précision ne peut être mesurée qu'à partir de données de référence spécifiques à la tâche, qui ne concernent pas les réécritures des DMP des patients, mais les DMP originaux des patients.

Dans ce travail, nous soulevons ces limites par une approche d'apprentissage joint de la réécriture des DMP et de la tâche de prédiction clinique détaillée dans la section suivante. La Figure 1 présente les différences entre notre approche et celles des travaux de l'état de l'art.

2 Description de l'approche

Dans cet article, nous proposons **Rewrite To Predict (ReToP)**, un modèle de prédiction clinique basé sur les LLMs. Comme montré sur la Figure 1 (approach c), **ReToP** est composé de deux modèles. Le premier est un LLM pour la réécriture du DMP et le second est un prédicteur de tâche clinique. Cette architecture est généralisable à différentes tâches prédictive avec un coût limité au réentraînement du prédicteur sans totalement réentraîner le modèle de réécriture.

Pour faire face à la haute dimensionnalité et à l'aspect "creux" des caractéristiques cliniques et être en mesure de générer des descriptions fidèles des DMP (défi **C1**), nous entraînons le module de réécriture des DMP à l'aide d'un ensemble de pseudo-étiquettes synthétiques de réécriture de patients en deux étapes. Nous échantillons d'abord les caractéristiques des DMP afin de construire des paraphrases diverses pour réduire la dimensionnalité tout en couvrant les caractéristiques saillantes des DMP originaux. Ensuite, nous sélectionnons les K meilleurs DMP paraphrasés comme pseudo-étiquettes en fonction d'un score de pertinence qui quantifie leur utilité clinique pour un ensemble de tâches cliniques représentatives. Les pseudo-étiquettes sélectionnées sont ensuite utilisées pour affiner le module de réécriture du DMP basé sur le LLM. Afin d'aligner la qualité de la réécriture des DMP des patients sur la précision des prédictions cliniques (défi **C2**), nous proposons un entraînement supervisé de bout en bout du modèle de réécriture de DMP et du prédicteur clinique. Le prédicteur clinique fournit une supervision itérative au modèle de réécriture de DMP basé sur le LLM, qui peut alors optimiser ses paramètres afin de générer des réécritures de patients qui facilitent la prédiction. La supervision du prédicteur quantifie l'impact de chaque réécriture de patient sur la précision de la prédiction en utilisant un score CSC (Classifier Supervised Contribution) inspiré de la mesure LSR (LM Supervised Retriever) (Shi *et al.*, 2024).

Le modèle d'apprentissage joint de réécriture-prédiction est présenté sur la Figure 2.

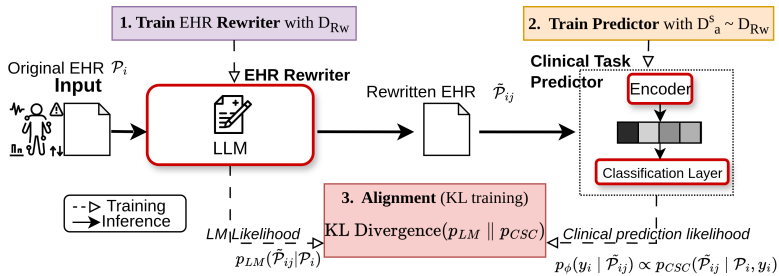


FIGURE 2 – Aperçu de l’entraînement et de l’inférence de ReToP : entraînement du modèle de réécriture avec les réécritures synthétiques D_{Rw} (1), puis entraînement conjoint avec le prédicteur sur la tâche cible (2–3).

Modèle	MIMIC-IV					
	MOR		RA		LOS	
	ROC(\uparrow)	PRC(\uparrow)	ROC(\uparrow)	PRC(\uparrow)	ROC(\uparrow)	PRC(\uparrow)
<i>EHR-oriented models</i>						
(a) RNN (Cho <i>et al.</i> , 2014))	63.33 \pm 3.4	3.38 \pm 0.7	65.76 \pm 1.2	72.49 \pm 1.3	73.02 \pm 0.7	58.01 \pm 1.3
(b) RETAIN (Choi <i>et al.</i> , 2016)	59.13 \pm 3.9	4.06 \pm 1.8	64.57 \pm 1.2	70.74 \pm 1.4	72.91 \pm 0.7	58.21 \pm 1.3
(c) GRASP (Zhang <i>et al.</i> , 2021)	58.94 \pm 3.6	2.68 \pm 0.5	63.95 \pm 1.2	68.92 \pm 1.4	70.48 \pm 0.8	53.85 \pm 1.3
(d) Llemr (Wu <i>et al.</i> , 2024)	52.07 \pm 4.1	3.17 \pm 1.3	59.20 \pm 1.1	64.89 \pm 1.3	75.36 \pm 0.7	62.21 \pm 1.3
(e) REmed (Kim <i>et al.</i> , 2024)	52.40 \pm 3.0	2.16 \pm 0.5	68.26 \pm 1.3	75.31 \pm 1.9	80.15 \pm 1.9	68.06 \pm 2.9
<i>Serialized Classifiers</i>						
(f) BioMedBERT (Gu <i>et al.</i> , 2021)	64.24 \pm 3.6	3.38 \pm 0.7	70.36 \pm 1.0	76.66 \pm 1.0	77.59 \pm 0.7	64.42 \pm 1.3
(g) Qwen (Yang <i>et al.</i> , 2024)	53.87 \pm 4.1	3.12 \pm 0.9	61.06 \pm 1.1	66.15 \pm 1.3	73.78 \pm 0.7	60.75 \pm 1.3
(h) Llama (Dubey <i>et al.</i> , 2024)	59.45 \pm 4.0	2.86 \pm 0.6	64.93 \pm 1.1	70.39 \pm 1.3	75.09 \pm 0.7	62.42 \pm 1.2
(i) ModernBERT (Warner <i>et al.</i> , 2024)	58.01 \pm 3.7	3.26 \pm 1.2	70.98 \pm 1.0	77.25 \pm 1.0	80.40 \pm 0.6	68.42 \pm 1.2
<i>Rewrite-then-Predict models</i>						
(j) self-gen (Sui <i>et al.</i> , 2024)	58.38 \pm 3.6	2.95 \pm 0.9	69.19 \pm 1.0	76.24 \pm 1.0	79.02 \pm 0.6	67.04 \pm 1.2
(k) π_t^h	61.00 \pm 4.2	3.58 \pm 1.0	65.25 \pm 1.0	73.77 \pm 1.1	70.04 \pm 0.7	54.74 \pm 1.3
(l) π_v^h	56.05 \pm 3.7	2.57 \pm 0.5	64.47 \pm 1.0	70.06 \pm 1.3	66.63 \pm 0.8	50.79 \pm 1.3
(m) π_{mrmr}^d (Lewis, 1992)	56.39 \pm 3.8	2.93 \pm 0.7	67.98 \pm 1.0	74.57 \pm 1.1	76.13 \pm 0.7	60.57 \pm 1.2
(n) π_{mrmr}^d (Ding & Peng, 2005)	60.00 \pm 3.6	3.21 \pm 0.8	63.08 \pm 1.1	68.10 \pm 1.3	79.48 \pm 0.6	68.53 \pm 1.2
(o) π_{rfe}^d (Guyon <i>et al.</i> , 2002)	62.54 \pm 3.5	3.69 \pm 1.3	62.03 \pm 1.0	67.55 \pm 1.3	76.84 \pm 0.7	63.22 \pm 1.2
(p) π_v^f	55.82 \pm 3.7	2.63 \pm 0.7	66.60 \pm 1.0	72.34 \pm 1.2	73.33 \pm 0.7	59.29 \pm 1.3
(q) π_v^f	59.37 \pm 3.8	3.73 \pm 1.5	68.03 \pm 1.0	74.37 \pm 1.1	77.51 \pm 0.7	63.41 \pm 1.2
ReToP _{Qwen} (ours)	69.75 \pm 3.8	4.61 \pm 1.0	71.45 \pm 0.9	77.92 \pm 1.0	80.38 \pm 0.6	68.91 \pm 1.2
ReToP _{Llama} (ours)	71.92 \pm 3.6	6.20 \pm 2.2	72.05 \pm 0.9	78.25 \pm 1.0	80.62 \pm 0.6	69.05 \pm 1.2

TABLE 1 – AUC-ROC (ROC) et AUC-PRC (PRC) scores sur MIMIC-IV. **Gras** = meilleur score, et le souligné = 2e meilleur score.

Modèle	MIMIC-IV →eICU			→EFEMERIS
	MOR	RA	LOS	MALF
ModernBERT	81.30/21.86	86.71/23.80	93.19/56.67	49.79/6.44
ReToP _{Llama}	81.42/22.29	86.72/23.78	93.16/56.78	51.53/6.93
ReToP _{Qwen}	81.40/21.93	86.79/23.91	93.17/ 56.80	50.41/ 7.08

TABLE 2 – Performances de généralisabilité avec les mesures AUC-ROC/AUC-PRC.

3 Expérimentations et Principaux Résultats

Nous avons mené des expérimentations sur plusieurs corpus de données DMP, notamment MIMIC-IV (Johnson *et al.*, 2023), eICU (Pollard *et al.*, 2018) et EFEMERIS (I *et al.*, 2009), ainsi qu’une série de tâches de prédiction clinique représentatives : prédiction de la mortalité, prédiction des réadmissions, prédiction de la durée du séjour et prédiction des malformations congénitales. Les résultats montrent que le cadre ReToP surpasse les modèles de l’état de l’art les plus récents pour ces tâches cliniques, avec des améliorations pouvant atteindre 23% du score AUC-ROC et un effet positif de chacun de ses composants (voir Table 2).

De plus, nous montrons que notre cadre est généralisable à des données et des tâches hors domaine, collections eICU (Pollard *et al.*, 2018) et EFEMERIS (I *et al.*, 2009), avec un coût de calcul faible lors du réglage fin du prédicteur clinique, sans réglage fin supplémentaire du modèle de réécriture du DMP (voir Table 2).

4 Conclusion

Nous avons présenté **ReToP**, un nouveau modèle qui exploite les LLM pour améliorer les performances de prédiction clinique. **ReToP** surpasse de façon significative un ensemble représentatif de modèles de base dans des tâches de prédiction clinique. Le cadre que nous proposons présente une capacité de transfert raisonnable vers des ensembles de données et des tâches hors domaine. Les travaux futurs pourraient explorer l’extension de ce cadre à un éventail plus large de tâches cliniques, y compris les tâches de classification multi-étiquettes, et assurer le juste compromis entre les performances du modèle l’ explicabilité de ses prédictions pour les experts du domaine.

Références

CHEN J., LIN H., HAN X. & SUN L. (2024). Benchmarking large language models in retrieval-augmented generation. *Proceedings of the AAAI Conference on Artificial Intelligence*, **38**(16), 17754–17762. DOI : [10.1609/aaai.v38i16.29728](https://doi.org/10.1609/aaai.v38i16.29728).

CHO K., VAN MERRIËNBOER B., BAHDANAU D. & BENGIO Y. (2014). On the properties of neural machine translation : Encoder–decoder approaches. In D. WU, M. CARPUAT, X. CARRERAS & E. M. VECCHI, Éd.s., *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, p. 103–111, Doha, Qatar : Association for Computational Linguistics. DOI : [10.3115/v1/W14-4012](https://doi.org/10.3115/v1/W14-4012).

- CHOI E., BAHADORI M. T., SUN J., KULAS J., SCHUETZ A. & STEWART W. (2016). Retain : An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems*, **29**.
- CHOI E., XIAO C., STEWART W. F. & SUN J. (2018). Mime : multilevel medical embedding of electronic health records for predictive healthcare. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, p. 4552–4562, Red Hook, NY, USA : Curran Associates Inc.
- CHOI E., XU Z., LI Y., DUSENBERRY M., FLORES G., XUE E. & DAI A. (2020). Learning the graphical structure of electronic health records with graph convolutional transformer. *Proceedings of the AAAI Conference on Artificial Intelligence*, **34**(01), 606–613. DOI : [10.1609/aaai.v34i01.5400](https://doi.org/10.1609/aaai.v34i01.5400).
- DING C. & PENG H. (2005). Minimum redundancy feature selection from microarray gene expression data. *Journal of bioinformatics and computational biology*, **3**(02), 185–205.
- DLIGACH D. & MILLER T. (2018). Learning patient representations from text. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, p. 119–123.
- DUBEY A., JAUHRI A., PANDEY A., KADIAN A., AL-DAHLE A., LETMAN A., MATHUR A., SCHELLEN A., YANG A., FAN A. *et al.* (2024). The llama 3 herd of models. *arXiv e-prints*, p. arXiv-2407.
- GU Y., TINN R., CHENG H., LUCAS M., USUYAMA N., LIU X., NAUMANN T., GAO J. & POON H. (2021). Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans. Comput. Healthcare*, **3**(1). DOI : [10.1145/3458754](https://doi.org/10.1145/3458754).
- GUYON I., WESTON J., BARNHILL S. & VAPNIK V. (2002). Gene selection for cancer classification using support vector machines. *Machine learning*, **46**(1), 389–422.
- HUR K., OH J., KIM J., KIM J., LEE M. J., CHO E., MOON S.-E., KIM Y.-H., ATALLAH L. & CHOI E. (2023). Genhpf : General healthcare predictive framework for multi-task multi-source learning. *IEEE Journal of Biomedical and Health Informatics*, **28**(1), 502–513.
- IL., C H., M F S., C G., A B., M G., C A.-C., R B., E E., J L M. & C D.-M. (2009). Prescription of drugs during pregnancy : a study using efemeris, the new french database. *European journal of clinical pharmacology*, **65**(8), 839–846.
- JIANG P., XIAO C., CROSS A. R. & SUN J. (2024). Graphcare : Enhancing healthcare predictions with personalized knowledge graphs. In *The Twelfth International Conference on Learning Representations*.
- JOHNSON A. E., BULGARELLI L., SHEN L., GAYLES A., SHAMMOUT A., HORNG S., POLLARD T. J., HAO S., MOODY B., GOW B. *et al.* (2023). MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data*, **10**(1), 1.
- K S., R Z., P J. & N. M. P. (2022). A systematic review of the prediction of hospital length of stay : Towards a unified framework. *PLOS Digit Health*, **14**. DOI : [10.1371/journal.pdig.0000017](https://doi.org/10.1371/journal.pdig.0000017).
- KIM J., SHIM C., YANG B. S. K., IM C., LIM S. Y., JEONG H.-G. & CHOI E. (2024). General-purpose retrieval-enhanced medical prediction model using near-infinite history. In K. DESHPANDE, M. FITERAU, S. JOSHI, Z. LIPTON, R. RANGANATH & I. URTEAGA, Éd.s., *Proceedings of the 9th Machine Learning for Healthcare Conference*, volume 252 de *Proceedings of Machine Learning Research* : PMLR.
- KWEON S., KIM J., KIM J., IM S., CHO E., BAE S., OH J., LEE G., MOON J. H., YOU S. C., BAEK S., HAN C. H., JUNG Y. B., JO Y. & CHOI E. (2024). Publicly shareable clinical large language model built on synthetic clinical notes. In L.-W. KU, A. MARTINS & V. SRIKUMAR, Éd.s., *Findings of the Association for Computational Linguistics : ACL 2024*, p. 5148–5168, Bangkok, Thailand : Association for Computational Linguistics. DOI : [10.18653/v1/2024.findings-acl.305](https://doi.org/10.18653/v1/2024.findings-acl.305).

LANDI I., GLICKSBERG B. & LEE H. E. A. (2020). Deep representation learning of electronic health records to unlock patient stratification at scale. *npj Digit. Med.*, **38**(3).

LEWIS D. D. (1992). Feature selection and feature extraction for text categorization. In *Speech and Natural Language : Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*.

LI R., WANG X. & YU H. (2024). LlamaCare : An instruction fine-tuned large language model for clinical NLP. In N. CALZOLARI, M.-Y. KAN, V. HOSTE, A. LENCI, S. SAKTI & N. XUE, Éd.s., *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, p. 10632–10641, Torino, Italia : ELRA and ICCL.

LIU H., WANG S., ZHU Y., DONG Y. & LI J. (2024a). Knowledge graph-enhanced large language models via path selection. In L.-W. KU, A. MARTINS & V. SRIKUMAR, Éd.s., *Findings of the Association for Computational Linguistics : ACL 2024*, p. 6311–6321, Bangkok, Thailand : Association for Computational Linguistics. DOI : [10.18653/v1/2024.findings-acl.376](https://doi.org/10.18653/v1/2024.findings-acl.376).

LIU Q., WU X., ZHAO X., ZHU Y., XU D., TIAN F. & ZHENG Y. (2024b). When moe meets llms : Parameter efficient fine-tuning for multi-task medical applications. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '24*, p. 1104–1114, New York, NY, USA : Association for Computing Machinery. DOI : [10.1145/3626772.3657722](https://doi.org/10.1145/3626772.3657722).

LOVÓN-MELGAREJO J., MORENO J. G., DAMASE-MICHEL C. & TAMINE L. (2026). Retop : Learning to rewrite electronic health records for clinical prediction. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining*.

LOVÓN-MELGAREJO J., MOUYSET M., OLEIWAN J., MORENO J. G., DAMASE-MICHEL C. & TAMINE L. (2025). Evaluating LLM abilities to understand tabular electronic health records : A comprehensive study of patient data extraction and retrieval. In *European Conference in Information Retrieval (ECIR)* : Elsevier.

MADHUMITA S., SIMON U., KIM L. & WALTER D. (2018). Patient representation learning and interpretable evaluation using clinical notes. *Journal of Biomedical Informatics*, **84**, 103–113.

NAIK A., PARASA S., FELDMAN S., WANG L. L. & HOPE T. (2022). Literature-augmented clinical outcome prediction. In M. CARPUAT, M.-C. DE MARNEFFE & I. V. MEZA RUIZ, Éd.s., *Findings of the Association for Computational Linguistics : NAACL 2022*, p. 438–453, Seattle, United States : Association for Computational Linguistics.

PETRONI F., ROCKTÄSCHEL T., RIEDEL S., LEWIS P., BAKHTIN A., WU Y. & MILLER A. (2019). Language models as knowledge bases? In K. INUI, J. JIANG, V. NG & X. WAN, Éd.s., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, p. 2463–2473, Hong Kong, China : Association for Computational Linguistics. DOI : [10.18653/v1/D19-1250](https://doi.org/10.18653/v1/D19-1250).

POLLARD T. J., JOHNSON A. E., RAFFA J. D., CELI L. A., MARK R. G. & BADAWI O. (2018). The eicu collaborative research database, a freely available multi-center database for critical care research. *Scientific data*, **5**(1), 1–13.

SHI W., MIN S., YASUNAGA M., SEO M., JAMES R., LEWIS M., ZETTLEMOYER L. & YIH W.-T. (2024). Replug : Retrieval-augmented black-box language models. In *NAACL-HLT*.

SINGHA A., CAMBRONERO J., GULWANI S., LE V. & PARNIN C. (2023). Tabular representation, noisy operators, and impacts on table structure understanding tasks in llms. In *Table Representation Learning Workshop at NeurIPS 2023*.

SUI Y., ZHOU M., ZHOU M., HAN S. & ZHANG D. (2024). Table meets llm : Can large language models understand structured table data ? a benchmark and empirical study. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining, WSDM '24*, p. 645–654, New York, NY, USA : Association for Computing Machinery. DOI : [10.1145/3616855.3635752](https://doi.org/10.1145/3616855.3635752).

TAMINE L. & GOEURIOT L. (2021). Semantic information retrieval on medical texts : Research challenges, survey, and open issues. *ACM Comput. Surv.*, **54**(7). DOI : [10.1145/3462476](https://doi.org/10.1145/3462476).

WARNER B., CHAFFIN A., CLAVIÉ B., WELLER O., HALLSTRÖM O., TAGHADOUINI S., GALLAGHER A., BISWAS R., LADHAK F., AARSEN T., COOPER N., ADAMS G., HOWARD J. & POLI I. (2024). Smarter, better, faster, longer : A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference.

WU Z., DADU A., NALLS M., FAGHRI F. & SUN J. (2024). Instruction tuning large language models to understand electronic health records. In A. GLOBERSON, L. MACKEY, D. BELGRAVE, A. FAN, U. PAQUET, J. TOMCZAK & C. ZHANG, Éd.s., *Advances in Neural Information Processing Systems*, volume 37, p. 54772–54786 : Curran Associates, Inc.

XU R., SHI W., YU Y., ZHUANG Y., JIN B., WANG M. D., HO J. & YANG C. (2024). Ram-ehr : Retrieval augmentation meets clinical predictions on electronic health records. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2 : Short Papers)*, p. 754–765.

YANG A., YANG B., HUI B., ZHENG B., YU B., ZHOU C., LI C., LI C., LIU D., HUANG F., DONG G., WEI H., LIN H., TANG J., WANG J., YANG J., TU J., ZHANG J., MA J., XU J., ZHOU J., BAI J., HE J., LIN J., DANG K., LU K., CHEN K., YANG K., LI M., XUE M., NI N., ZHANG P., WANG P., PENG R., MEN R., GAO R., LIN R., WANG S., BAI S., TAN S., ZHU T., LI T., LIU T., GE W., DENG X., ZHOU X., REN X., ZHANG X., WEI X., REN X., FAN Y., YAO Y., ZHANG Y., WAN Y., CHU Y., LIU Y., CUI Z., ZHANG Z. & FAN Z. (2024). Qwen2 technical report. *arXiv preprint arXiv :2407.10671*.

ZHANG C., GAO X., MA L., WANG Y., WANG J. & TANG W. (2021). Grasp : generic framework for health status representation learning based on incorporating knowledge from similar patients. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, p. 715–723.

ZHU Y., WANG Z., HE L., XIE S., ZHENG X., MA L. & PAN C. (2024). Prism : Mitigating ehr data sparsity via learning from missing feature calibrated prototype patient representations. *CIKM '24*, p. 3560–3569, New York, NY, USA : Association for Computing Machinery. DOI : [10.1145/3627673.3679521](https://doi.org/10.1145/3627673.3679521).



PDF Download
3773966.3777969.pdf
11 March 2026
Total Citations: 0
Total Downloads: 165

Latest updates: <https://dl.acm.org/doi/10.1145/3773966.3777969>

RESEARCH-ARTICLE

ReToP: Learning to Rewrite Electronic Health Records for Clinical Prediction

JESUS LOVON-MELGAREJO, University of Toulouse, Toulouse, Occitanie, France

JOSE G. MORENO, University of Toulouse, Toulouse, Occitanie, France

CHRISTINE DAMASE-MICHEL, Toulouse University Hospital, Toulouse, Occitanie, France

LYNDA TAMINE, University of Toulouse, Toulouse, Occitanie, France

Open Access Support provided by:

Toulouse University Hospital

University of Toulouse

Published: 21 February 2026

Citation in BibTeX format

WSDM '26: The Nineteenth ACM International Conference on Web Search and Data Mining
February 22 - 26, 2026
ID, Boise, USA

Conference Sponsors:

SIGKDD
SIGWEB
SIGIR
SIGMOD

ReToP: Learning to Rewrite Electronic Health Records for Clinical Prediction

Jesus Lovon-Melgarejo
jesus.lovon@irit.fr
University of Toulouse, IRIT
Toulouse, France

Christine Damase-Michel
christine.damase-michel@outoulouse.fr
Toulouse University Hospital
Toulouse, France
University of Toulouse, Inserm UMR 1295
CERPOP-SPHERE Team
Toulouse, France

Jose G. Moreno
Jose.Moreno@irit.fr
University of Toulouse, IRIT
Toulouse, France

Lynda Tamine
Lynda.Tamine@irit.fr
University of Toulouse, IRIT
Toulouse, France

Abstract

Electronic Health Records (EHRs) provide crucial information for clinical decision-making. However, their high-dimensionality, heterogeneity, and sparsity make clinical prediction challenging. Large Language Models (LLMs) allowed progress towards addressing this challenge by leveraging parametric medical knowledge to enhance EHR data for clinical prediction tasks. Despite the significant achievements made so far, most of the existing approaches are fundamentally task-agnostic in the sense that they deploy LLMs as EHR encoders or EHR completion modules without fully integrating signals from the prediction tasks. This naturally hinders task performance accuracy. In this work, we propose **Rewrite-To-Predict (ReToP)**, an LLM-based framework that addresses this limitation through an end-to-end training of an EHR rewriter and a clinical predictor. To cope with the lack of EHR rewrite training data, we generate synthetic pseudo-labels using clinical-driven feature selection strategies to create diverse patient rewrites for fine-tuning the EHR rewriter. **ReToP** aligns the rewriter with prediction objectives using a novel *Classifier Supervised Contribution (CSC)* score that enables the EHR rewriter to generate clinically relevant rewrites that directly enhance prediction. Our **ReToP** framework surpasses strong baseline models across three clinical tasks on MIMIC-IV. Moreover, the analysis of **ReToP** shows its generalizability to unseen datasets and tasks with minimal fine-tuning while preserving faithful rewrites and emphasizing task-relevant predictive features.

CCS Concepts

• Computing methodologies → Natural language processing.

Keywords

LLMs, Electronic Health Record (EHR), Clinical Prediction

ACM Reference Format:

Jesus Lovon-Melgarejo, Jose G. Moreno, Christine Damase-Michel, and Lynda Tamine. 2026. ReToP: Learning to Rewrite Electronic Health Records for Clinical Prediction. In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining (WSDM '26)*, February 22–26, 2026, Boise, ID, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3773966.3777969>

1 Introduction

Clinical predictive models heavily rely on electronic health records (EHRs), which encode longitudinal patients' medical features (e.g., disease, procedures) to estimate the risks of having or developing a health-related outcome. Clinical prediction faces several challenges, among which are heterogeneity, high dimensionality, and sparsity of EHRs [24, 58]. Significant research effort has been dedicated to tackle these challenges through the development of machine learning models for a wide-range of clinical prediction tasks including diagnosis prediction [8, 15, 22, 34], mortality prediction [5, 15, 34, 36], readmission prediction [15, 18] and length of stay prediction [15, 21]. One major agreed-upon finding from these studies is that leveraging expert knowledge with patient data insights enhances EHR modeling, which leads to significant improvement in prediction accuracy. Previous studies mostly integrated expert symbolic knowledge in the form of clinical knowledge graphs that represent medical features (e.g., ICD10 codes of diagnoses) and their relationships to model valuable contextual information to complement EHR data [4, 5, 47, 51].

Subsequently, Large Language Models (LLMs) have emerged as a significant milestone in the transition from symbolic knowledge to parametric knowledge, bringing us one step closer to potential knowledge graphs with impressive capabilities of language understanding and generation [38]. To date, there is a large body of work showing the potential of LLMs to perform medical tasks [23, 28, 49]. In the literature related to LLM-based clinical prediction, models rely on either of the following approaches illustrated in Figure 1: (a) leveraging fine-tuned pre-trained LLM specialized on medical textual corpora (e.g., Llamacare [28], Llmr [49]). Although showing strong capabilities in interpreting complex EHR features, these works have been evaluated particularly on medical knowledge understanding tasks (e.g., medical question-answering) and



This work is licensed under a Creative Commons Attribution 4.0 International License.
WSDM '26, Boise, ID, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2292-9/2026/02
<https://doi.org/10.1145/3773966.3777969>

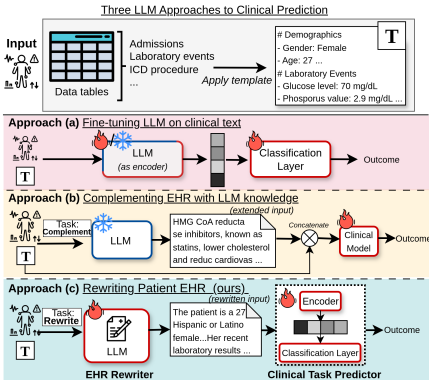


Figure 1: Comparison of three LLM-based approaches for clinical prediction from EHRs. Approach (a) relies on fine-tuned LLMs as EHR encoders. Approach (b) relies on LLMs to complement EHRs. Our approach, ReToP (c), uses a trainable LLM-based EHR rewriter aligned with clinical tasks.

usually require a significant fine-tuning cost for achieving generalization ability [32]; (b) complementing EHRs with clinical knowledge from LLMs either agnostically [51] or dependently of EHR data before training the predictive model [18]. However, these existing works have two main limitations: (1) they have low flexibility in determining to what extent a certain knowledge is useful vs. noisy when naively fusing between LLM parametric knowledge and KG knowledge [1, 31]; (2) they do not link EHR completion with clinical prediction. Since knowledge graphs contain a large amount of relational information, complementing EHR with noisy data may hinder the prediction accuracy.

Inspired by how human experts diagnose, we raise a critical question: “Given the patient EHR, what are the clinical rationales to support prediction?” Our work answers this question while tackling the above-cited limitations. We design an LLM-based EHR rewriter, acting as the unique source of expert knowledge, able to generate faithful and salient descriptions of the original patient EHR. The LLM-based EHR rewriter is further fine-tuned to revise the EHR rewrites to better support the clinical prediction. However, the implementation of this answer poses the two following challenges:

- **C1:** Given an input patient EHR, the LLM-based rewriter might be able to faithfully recall among a huge number of clinical features, the most salient ones, for any prediction task. A potential solution is using an off-the-shelf LLM-based EHR rewriter. However, previous work showed that LLMs still struggle to comprehend tabular data, including EHRs [33, 45, 46]. Since the gold relevant clinical features of an EHR are a priori unknown, the alternative of designing a trainable LLM-based EHR rewriter poses the challenge of a lack of gold supervision.

- **C2:** The relevant clinical features of EHRs are likely to be task-dependent. Thus, prediction accuracy might serve as a proxy evaluation of the quality of the EHR rewrite. However, accuracy can only be measured on gold task-specific data, which do not involve patient EHR rewrites but original patient EHR instead.

In this work, we introduce **Rewrite To Predict (ReToP)**, a new LLM-based framework for clinical prediction. As shown in Figure 1 (approach c), **ReToP** is composed of two modules. The first module represents an LLM-based EHR rewriter, and the second module is a clinical task predictor. Our framework is generalizable to different predictive clinical tasks at the light cost of fine-tuning the predictor without fully re-training the LLM-based rewriter.

To cope with the high dimensionality and sparsity of clinical features and be able to generate faithful EHR descriptions (challenge C1), we train the EHR rewriter module using a set of synthetic pseudo-labels of patient-rewrites in two stages. We first sample EHR features to build diverse paraphrases to reduce dimensionality while covering salient features of the original EHR. Then, we select the top-K paraphrased EHRs as pseudo-labels based on a relevance score that quantifies their clinical utility to a set of representative clinical tasks. The selected pseudo-labels are then used for fine-tuning the LLM-based EHR rewriter. To align patient EHR rewrite quality to clinical prediction accuracy (challenge C2), we propose an end-to-end supervised training of the EHR rewriter and clinical predictor. The clinical predictor iteratively provides supervision to the LLM-based EHR rewriter, which can then optimize its parameters to generate patient rewrites that help prediction. The predictor supervision quantifies the impact of each patient rewrite on prediction accuracy by using a Classifier Supervised Contribution (CSC) score inspired by the LM Supervised Retriever (LSR) measure [43].

We conduct extensive experiments across multiple EHR datasets, including MIMIC-IV [20], eICU [39], and EFEMERIS [16], as well as a set of representative clinical prediction tasks: mortality prediction, readmission prediction, length of stay prediction, and congenital malformation prediction. The results show that the **ReToP** framework outperforms recent state-of-the-art baselines across these clinical tasks, achieving improvements up to 23% in AUC-ROC score with a positive effect of each of its components. Furthermore, we show that our framework is generalizable to out-of-domain data and tasks with light computational cost during fine-tuning of the clinical predictor, without additional fine-tuning of the LLM-based EHR rewriter. Finally, qualitative analysis reveals that our EHR rewriter preserves faithfulness to the original EHR, and that KL alignment can emphasize serendipitous clinical features, not expected by experts, although valuable for accurate prediction.

2 Related Work

2.1 Clinical prediction

Clinical decision-making has made significant progress since deploying deep learning models on EHRs [50]. Early models have particularly tackled the challenge of domain-specific adaptation and demonstrated their ability to improve the effectiveness of a wide range of predictive clinical tasks, including health failure prediction [3, 4, 55], mortality prediction [5, 15, 34, 36], and diagnosis prediction [8, 15, 22, 34, 36].

The follow-up rise of pre-trained language models (LMs) has led to a vast amount of research investigating their capabilities to tackle the challenge of EHR data scarcity and clinical language understanding tasks. Among these models, BEHRT [29], MedBERT [40], TransformEHR [54] and ClinicalBERT [14] have particularly led to progress in disease prediction. Recently, REMED [22] proposed a retrieval-based filtering approach to address the sequence length limitations of LMs when processing EHR events.

In the same line of transformer-based models, but with more impressive transfer capabilities, decoder-only LLMs have emerged as promising tools for predictive healthcare. The first category of work focused on the design of foundation models on EHRs to tailor their inherent abilities of multi-tasking and reasoning to clinical applications [23, 28, 49]. For instance, Llemr [49] is an instruction-tuned multimodal model based on Llama [30] that jointly encodes vectorized EHR event representations with natural language questions, evaluated on EHR-based question-answering tasks, and adapted for clinical prediction tasks, including mortality, readmission, length-of-stay, and diagnosis prediction. The second category of work integrates LLMs in the pipeline of a clinical predictive task by achieving various roles such as supporting clinical reasoning [37, 44] or completing EHR descriptions to improve the prediction [18, 32, 37, 51]. Especially, Jiang et al. [18] propose a framework that relies on LLMs prompted to build a graph-based patient representation, which is then used to train a deep neural healthcare predictive model based on a graph neural network (GNN). Unlikely, Nguyen et al. [37] adopt a retrieval-augmented approach to enhance patient representations using multiple domain-specific resources. The authors leverage the LLMs to generate a local patient representation, a summary of knowledge resulting from a retriever queried with concepts from the EHR. This representation is then complemented by a patient-visit representation, and both are trained to achieve a predictive task.

Our work is significantly different in that we attempt to improve clinical prediction by fine-tuning, instead of zero-shot prompting, LLMs to generate the patient EHR rewrites without any external knowledge. Furthermore, instead of using task-agnostic enhanced EHR representations as input, we show how to train an LLM-based EHR rewriter using clinical prediction outcomes as supervision.

2.2 LLMs as feature selectors

Recent work has increasingly explored the application of LLMs as feature selectors [6, 17, 27], leveraging their multi-task ability to shift from traditional statistical methods toward more semantically-aware feature selection approaches. Early work [6] introduced a prompt-based approach that formulates feature selection as a binary classification problem. The models generate “Yes” or “No” responses based on log-probability differences between these tokens to indicate feature importance given task descriptions, target features, and feature descriptions. LLM-Select [17] extensively explored zero-shot and few-shot approaches across multiple paradigms, using score-based, rank-based, and dialogue-based methods to elicit feature relevance. Moreover, to address challenges in specialized domains where LLMs may have limited knowledge of domain-specific features, such as healthcare, RAFS [27] proposed a retrieval-augmented

generation (RAG) approach that retrieves relevant feature descriptions from external knowledge sources.

However, these approaches mainly rely on feature descriptions, which require specialized domain knowledge that becomes resource-intensive when handling large feature sets. Additionally, few-shot approaches remain highly susceptible to prompt variations [42, 56]. In contrast, our work leverages the LLM’s internal knowledge through controlled generation by fine-tuning an LLM-based EHR rewriter that implicitly operates on feature importance derived from statistical feature selection heuristics applied to raw EHR descriptions, rather than relying on EHR feature descriptions.

3 Background and Notations

An EHR is an individual patient record of a sequence of visits. In this section, we introduce the notions and notations used in our work and formulate the clinical prediction task.

3.1 Basic notions

- **Patient EHR.** A patient’s EHR P_i is composed of demographic information and longitudinal representation of the health conditions recorded during a sequence of hospital visits $V^1 \dots V^{|N_i|}$. Each visit V^t includes a set of feature-value tuples $\mathcal{T}_i^t = \{(f_k^t, val_k^t)\}_{k=1}^n$ where $t = \{1..T\}$ are discrete timestamps during the visit V^t , and feature $f_k \in \mathcal{F}$ with $\mathcal{F} = \{f_1, \dots, f_n\}$ a reference set of n features from different modalities (e.g., demographic, disease, medication).
- **Clinical prediction task.** We consider standard healthcare prediction tasks $s \in \mathbb{S}$, where the input space is a population of patient EHRs \mathbb{P} with each EHR $P_i \in \mathbb{P}$ in the form of tuples sampled from $\bigcup_{t=1}^T \mathcal{T}_i^t$ with $1..T$ discrete timestamps, and the output space is \mathcal{Y} with $y_i \in \mathcal{Y}$ a binary label, i.e., $y_i = \{0, 1\}$. Given a patient EHR P_i and a clinical task s , a predictive function f outputs the clinical outcome y_i . Without loss of generality, we consider the mortality prediction, readmission prediction, length of stay prediction, and congenital malformation prediction tasks (§5.2).
- **Patient EHR verbalisation.** As done in previous work [12, 33, 45, 46], a verbaliser function v is used to convert EHR feature-value tuples of a patient P_i (i.e., $P_i = \{\bigcup_{t=1}^T \mathcal{T}_i^t(t)\}$) into a natural language description \mathcal{P}_i . In this work, we applied markdown templates, verbalizing each tuple (f_k, val_k) into a string “- f_k ; val_k ”. We concatenate different feature modalities with headers “# {feature name}”.

To simplify the notation, we also refer to f as the predictive function that maps a text-based patient EHR $\mathcal{P}_i = v(P_i)$ to the clinical outcome y_i .

3.2 Problem statement and solution overview

ReToP employs two trainable parametrized models: a rewrite model M_θ and a task-specific predictor model f_ϕ . Formally, ReToP re-frames the prediction function f that supports the clinical task s , as a joint Rewrite (M_θ)-To-Predict (f_ϕ) model that computes likelihood of outcome y_i conditioned to patient EHR \mathcal{P}_i^s . To compute this likelihood, we adopt an ensemble strategy as done in previous

work [43]:

$$p(y_i^s | \mathcal{P}_i^s, R_w, \theta, \phi) = \sum_{\tilde{\mathcal{P}}_i^s \in R_w} p(y_i^s | \tilde{\mathcal{P}}_i^s, \theta, \phi) \times \delta(\tilde{\mathcal{P}}_i^s, \mathcal{P}_i^s) \quad (1)$$

where R_w includes all the possible rewrites of \mathcal{P}_i^s , limited in practice to n_i rewrites (§4.2.2), \oplus denotes the concatenation of two sequences, $p(y_i^s | \tilde{\mathcal{P}}_i^s, \theta, \phi)$ is the α -weighted linear combination computed by f_θ for outcome y_i^s conditioned on patient EHR \mathcal{P}_i^s and corresponding rewrite $\tilde{\mathcal{P}}_i^s$,

$$p(y_i^s | \tilde{\mathcal{P}}_i^s \oplus \tilde{\mathcal{P}}_i^s, \theta, \phi) = \alpha \times p(y_i^s | \tilde{\mathcal{P}}_i^s, \theta, \phi) + (1 - \alpha) \times p(y_i^s | \mathcal{P}_i^s, \theta, \phi). \quad (2)$$

$\delta(\tilde{\mathcal{P}}_i^s, \mathcal{P}_i^s)$, computed by M_θ , as the probability $M_\theta(\tilde{\mathcal{P}}_i^s | \mathcal{P}_i^s)$ of generating $\tilde{\mathcal{P}}_i^s$ as a rewrite of \mathcal{P}_i^s , θ and ϕ are model parameters trainable using task-specific training data $\mathcal{D} = \bigcup_{s \in \mathbb{S}} \mathcal{D}^s$ with $\mathcal{D}^s = \{(\mathcal{P}_i^s, y_i^s)\}_{1 \leq i \leq m_s}$ of m_s patients \mathcal{P}_i^s sampled from population $\mathbb{P}^s \subset \mathbb{P}$ with corresponding clinical outcomes y_i^s .

We leverage the capabilities of LLMs for text generation [26, 35] to support the patient rewriter model M_θ and use any backbone classifier model to support the prediction function f_ϕ .

To tackle challenge **C1**, we build a synthetic dataset \mathcal{D}_{R_w} across tasks to fine-tune M_θ . \mathcal{D}_{R_w} is composed of sets of m_i pseudo labels of rewrites to each seed patient \mathcal{P}_i , such as $\mathcal{D}_{R_w} = \{(\mathcal{P}_i, \tilde{\mathcal{P}}_{ij})\}_{\mathcal{P}_i \in \mathbb{P}, 1 \leq j \leq m_i}$. \mathcal{D}_{R_w} is built-up on a top-K selection strategy over candidate patient rewrites R_w generated using paraphrasing-based methods over original patient EHR \mathcal{P}_i .

To tackle challenge **C2**, **ReToP** is trained over a training task-specific dataset \mathcal{D}_s^* for end-to-end training of the patient rewriter M_θ and the prediction function f_ϕ of task s , such as: $\mathcal{D}_{P_r}^s = \{(\mathcal{P}_i^s, \tilde{\mathcal{P}}_{ij}^s, y_i^s)\}_{j=1 \dots n_i} | (\tilde{\mathcal{P}}_{ij}^s = M_\theta(\tilde{\mathcal{P}}_{ij}^s | \mathcal{P}_i^s) \wedge (\mathcal{P}_i^s, y_i^s) \in \mathcal{D}^s)$. Formally, the training of **ReToP** based on Eq. (1) can be reframed as an optimization problem where we seek among all the patient candidate rewrites $\tilde{\mathcal{P}}_{ij}^s$ generated from the original patient \mathcal{P}_i^s by model M_θ , the patient rewrite $\tilde{\mathcal{P}}_{ij}^s$ that maximizes the expectation of the clinical prediction task s as follows:

$$\tilde{\mathcal{P}}_{ij}^* = \arg \max_{\tilde{\mathcal{P}}_{ij}^s} \mathbb{E}_{\mathcal{D}_{R_w}} [M_\theta(\tilde{\mathcal{P}}_{ij}^s | \mathcal{P}_i^s)] = \arg \max_{\tilde{\mathcal{P}}_{ij}^s} \mathbb{E}_{\mathcal{D}_{P_r}^s} [f_\phi(y_i^s | \tilde{\mathcal{P}}_{ij}^s)] \quad (3)$$

4 Methodology

In this section, we describe the details of the procedure for building the synthetic training dataset (\mathcal{D}_{R_w}) used for fine-tuning the EHR rewriter (§4.1) and our end-to-end training procedure of the patient rewriter and the clinical prediction model (§4.2).

4.1 Fine-tuning the EHR rewriter

4.1.1 Synthetic training dataset generation. To enhance the LLM’s ability to rewrite patient EHRs for clinical tasks, we deploy a paraphrasing-based generation approach since it has been shown to favor diversity and faithfulness [19, 26]. Unlike previous work, we apply paraphrasing operators over the input patient EHR \mathcal{P}_i built upon feature-value tuples $\mathcal{T}_i^e = \{(f_k^e, val_k^e)\}_t$. We aim to reduce feature dimensionality and improve clinical feature diversity while remaining *faithful* to the original patient EHR. Specifically,

we consider a set of clinical tasks $s \in \mathbb{S}$ and we generate for each task and each seed patient P_i sampled from population $s\mathbb{P}^s \subset \mathbb{P}^s$, a set of K feature-based paraphrases $\{P_{ij}\}$ $j = 1 \dots K$ using operators dedicated to EHR data. In our work, we design $K = 8$ feature-based paraphrasing operators, inspired by previous work in prompt learning [27, 57], grouped by their feature selection strategy:

Heuristic-based π^h : we select clinical features based either on the temporal or value criteria given the importance of such criteria agnostic to task-specific outcome: 1) π_1^h : selects the $x\%$ more recent feature-value tuples such as $P_{i1} = \bigcup_{t=s}^{T_i} \mathcal{T}_i^e(t)$ with $s = (1-x) * T_i$; 2) π_2^h selects clinical features with abnormal values based on reference ranges such as $P_{i2} = \{(f_k^e, val_k^e)\}_t \in \mathcal{T}_i^e | (val_k^e < Min_k) \vee (val_k^e > Max_k)\}$ with $[Min_k \dots Max_k]$ is the range reference of feature f_k .

Data-driven π^d : we apply traditional feature selection methods that identify task-relevant features: 1) π_{mi}^d : filtering by mutual information [25], 2) π_{mrmr}^d : based on minimum redundancy maximum relevance selection [7], and 3) π_{frc}^d : recursive feature elimination [11]. Each method computes a score per feature, and we select the top $x\%$ with the highest scoring features. Formally, for $j \in \{3, 4, 5\}$ corresponding to methods $\{\pi_{mi}^d, \pi_{mrmr}^d, \pi_{frc}^d\}$ respectively, we define: $P_{ij} = \{(f_k^e, val_k^e)\}_t \in \mathcal{T}_i^e | \text{score}_j(f_k^e, \mathcal{D}^s) \geq x_k\}$, where $\text{score}_j(f_k^e, \mathcal{D}^s)$ denotes the score assigned to feature f_k by method j given the dataset \mathcal{D}^s , and x_k is the x^{th} percentile threshold of all feature scores for task s .

Random-based π^r : to favor the diversity of candidate clinical predictive features and values while reducing the dimensionality of the EHR, we design two operators: 1) π_r^r randomly selects a sample of clinical features and then corresponding feature-value tuples: $P_{i6} = \{(f_k^e, val_k^e)\}_t \in \mathcal{T}_i^e | f_k^e \in \mathcal{F}$; 2) π_o^r randomly selects a subset of feature-value tuples from P_i such as $P_{i7} = \{(f_k^e, val_k^e)\}_t \in \mathcal{T}_i^e$. Finally, we consider the trivial case of the identity operator 3) π_1^r that simply copies and pastes the original EHR: $P_{i8} = P_i$.

For each patient EHR P_i verbalized into $\{P_{ij}\}_{j=1 \dots 8}$ into a natural language patient description \mathcal{P}_{ij} using the verbalizer v . For each patient \mathcal{P}_i , we generate K input pairs $(\mathcal{P}_i, \mathcal{P}_{ij})$. Finally, we obtain for each task s the dataset $R_w^s = \{(\mathcal{P}_i, \tilde{\mathcal{P}}_{ij}^s)\}_{\mathcal{P}_i \in s\mathbb{P}^s}$.

4.1.2 Fine-tuning the EHR rewriter. The Algorithm 1 presents the pseudocode for building the synthetic training dataset \mathcal{D}_{R_w} and fine-tuning the patient rewriter M_θ given a set of clinical tasks \mathbb{S} . For each pair $(\mathcal{P}_i, \mathcal{P}_{ij}) \in R_w^s$, we evaluate the clinical relevance of the candidate rewrite $\tilde{\mathcal{P}}_{ij}$ for task s using a task-specific scorer $Score^s$. In our work, we consider $Score^s$ with the same architecture as our predictor model f_ϕ , trained on \mathcal{D}_{sub}^s , a subset of task-specific training data from \mathcal{D}^s involving patients with corresponding rewrites in R_w^s and their corresponding labels from the original EHR, such as $\mathcal{D}_{sub}^s = \{(\mathcal{P}_{ij}, y_i) | (\mathcal{P}_i, \mathcal{P}_{ij}) \in R_w^s, (\mathcal{P}_i, y_i) \in \mathcal{D}^s\}$. Our intuition is that $Score^s$ will assign higher predictive scores to informative rewrites while penalizing less relevant ones. For each patient \mathcal{P}_i , we select high-scoring rewrites as synthetic pseudo-labels into \mathcal{D}_{R_w} , as follows:

$$\mathcal{D}_{R_w} = \bigcup_{s \in \mathbb{S}} \{(\mathcal{P}_i, \tilde{\mathcal{P}}_{ij})_{1 \leq j \leq m_i} | Score^s(\tilde{\mathcal{P}}_{ij}) \geq \tau_s\}_{\mathcal{P}_i \in \mathcal{D}_{sub}^s} \quad (4)$$

Input: Task-specific rewrites $\{Rw^s\}$, scorer datasets $\{\mathcal{D}_{sub}^s\}$, base LLM rewriter M_θ^0 , selection percentile $k\%$, number of rewrites per patient K

Output: Fine-tuned EHR rewriter M_θ

$\mathcal{D}_{Rw} \leftarrow \emptyset;$

foreach $task\ s \in \mathcal{S}$ **do**

$Scorer^s \leftarrow \text{Train}(f_\phi, \mathcal{D}_{sub}^s)$ // Train scorer

$S^s \leftarrow \{Scorer^s(\mathcal{P}_{ij}) \mid (\mathcal{P}_i, \mathcal{P}_{ij}) \in Rw^s\}$ // Compute scores

$\tau_s \leftarrow \min\{x \in S^s : \text{rank}(x) \leq |S^s|k/100\}$ // Fix threshold

foreach $(\mathcal{P}_i, \mathcal{P}_{ij}) \in Rw^s$ **do**

$\hat{\mathcal{P}}_i^s \leftarrow \emptyset$ // Initialize candidate list

if $Scorer^s(\mathcal{P}_{ij}) \geq \tau_s$ **then**

$\hat{\mathcal{P}}_{ij} \leftarrow \mathcal{P}_{ij};$

$\mathcal{D}_{Rw} \leftarrow \mathcal{D}_{Rw} \cup \{(\mathcal{P}_i, \hat{\mathcal{P}}_i^s)\};$

end

end

end

$M_\theta \leftarrow \text{Fine-tune}(M_\theta^0, \mathcal{D}_{Rw});$

return $M_\theta;$

Algorithm 1: Fine-Tuning the EHR rewriter

with $m_i \leq K$ is the number of high-quality rewrites selected as pseudo labels for patient \mathcal{P}_i , τ_s is the score threshold obtained by $Scorer^s$ to keep the top- $k\%$ rewrites. As a result, we generate the unified dataset \mathcal{D}_{Rw} containing filtered rewrites across the set \mathcal{S} of clinical prediction tasks.

Finally, we fine-tune the patient EHR rewriter M_θ on \mathcal{D}_{Rw} using causal language modeling with a task-agnostic instruction.

4.2 End-to-End training of the EHR rewriter and the clinical predictor

In this section, we detail the calculation of the clinical prediction likelihood f_ϕ followed by the co-training methodology of the patient rewriter M_θ and the clinical predictor f_ϕ .

4.2.1 Computing the clinical prediction likelihood. The **ReToP** framework is trainable using task-specific data $\mathcal{D}^s = \{(\mathcal{P}_i^s, y_i^s)\}_{1 \leq i \leq m_s}$ where \mathcal{P}_i^s is sampled from population \mathbb{P}^s and $y_i^s \in \mathcal{Y}$ is the corresponding binary label. To enhance the clinical predictor within the rewrite-to-predict framework, we augment \mathcal{D}^s with training examples \mathcal{D}_a^s composed of rewrites from both paraphrase operators and our fine-tuned EHR rewriter, with all rewrites \mathcal{P}_{ij}^s inheriting labels y_i^s from their original EHR \mathcal{P}_i^s .

$$\mathcal{D}_a^s = \{(\mathcal{P}_{ij}^s, y_i^s) \mid (\mathcal{P}_i^s, \mathcal{P}_{ij}^s) \in Rw^s, (\mathcal{P}_i^s, y_i^s) \in \mathcal{D}^s\} \cup \{(M_\theta(\mathcal{P}_{ij}|\mathcal{P}_i), y_i^s) \mid (\mathcal{P}_i^s, y_i^s) \in \mathcal{D}^s\} \quad (5)$$

We design the clinical predictive function as an encoder classification model, optimized using Binary Cross-Entropy (BCE) loss:

$$\mathcal{L}_C = - \sum_{(\mathcal{P}_i^s, y_i^s) \in \mathcal{D}_a^s} [y_i \log p(y_i|\mathcal{P}_i, \phi) + (1 - y_i) \log(1 - p(y_i|\mathcal{P}_i, \phi))] \quad (6)$$

4.2.2 Alignment of the EHR rewriter with clinical prediction. As we formulate the overall clinical prediction task as an optimization problem of the EHR rewriter (§3.2), we propose a further fine-tuning of the EHR rewriter M_θ to align its generation with the likelihood of the clinical predictor f_ϕ , as presented in Figure 2. Specifically, we




Figure 2: Overview of ReToP training and inference. We train the EHR rewriter on \mathcal{D}_{Rw} (step 1). Then, jointly end-to-end train the rewriter and predictor on the target task (steps 2-3).

build a dual training dataset $\mathcal{D}_{Pr}^s = \{(\mathcal{P}_i^s, \hat{\mathcal{P}}_{ij}^s)_{j=1..m_i} \mid (\mathcal{P}_i^s, y_i^s) \in \mathcal{D}^s \wedge (\hat{\mathcal{P}}_{ij}^s = M_\theta(\mathcal{P}_{ij}|\mathcal{P}_i^s))\}$. This fine-tuning aims to align the patient EHR rewriter with the clinical prediction performance, resulting in more relevant patient rewrite generation. Inspired by previous work [41, 43], we introduce the *Classifier Supervised Tribution (CSC)* score, an adaptation of the LM-Supervised Retrieval (LSR) score [43], for clinical prediction tasks. Given a training triplet $(\mathcal{P}_i^s, \mathcal{P}_{ij}^s, y_i^s) \in \mathcal{D}_{Pr}^s$, the CSC score quantifies the relative effectiveness of patient rewrite $\hat{\mathcal{P}}_{ij}^s$ to help accurately predicting the target clinical outcome y_i^s regarding all the candidate rewrites $\{\hat{\mathcal{P}}_{ij}^s\}_{j=1}^{m_i}$ of patient \mathcal{P}_i^s :

$$p_{CSC}(\hat{\mathcal{P}}_{ij}^s \mid \mathcal{P}_i^s, y_i^s) = \frac{\exp(p(y_i^s \mid \hat{\mathcal{P}}_{ij}^s, \phi) / \tau)}{\sum_{(\mathcal{P}_{ij}^s, \hat{\mathcal{P}}_{ij}^s) \in \mathcal{D}_{Pr}^s} \exp(p(y_i^s \mid \hat{\mathcal{P}}_{ij}^s, \phi) / \tau)} \quad (7)$$

where τ is a temperature scaling parameter. The CSC score measures the contribution of a patient rewrite \mathcal{P}_{ij} to the correct classification outcome y_i^s . To make the patient rewriter model M_θ guided by the clinical predictor f_ϕ we minimize the Kullback-Leibler (KL) divergence between the language model’s output distribution and the CSC-weighted probability:

$$\mathcal{L}_{KL} = \mathbb{E}_D KL(p_{LM}(\hat{\mathcal{P}}_{ij} \mid \mathcal{P}_i^s, \theta) \parallel p_{CSC}(\hat{\mathcal{P}}_{ij} \mid \mathcal{P}_i^s, y_i^s)) \quad (8)$$

where $p_{LM}(\hat{\mathcal{P}}_{ij} \mid \mathcal{P}_i^s, \theta)$ represents the rewriter’s probability distribution over candidate rewrites, defined as:

$$p_{LM}(\hat{\mathcal{P}}_{ij} \mid \mathcal{P}_i^s, \theta) = \frac{\exp(p(\hat{\mathcal{P}}_{ij} \mid \mathcal{P}_i^s, \theta) / \kappa)}{\sum_{(\mathcal{P}_{ij}^s, \hat{\mathcal{P}}_{ij}^s) \in \mathcal{D}_{Pr}^s} \exp(p(\hat{\mathcal{P}}_{ij} \mid \mathcal{P}_i^s, \theta) / \kappa)} \quad (9)$$

with κ as a temperature scaling factor. The final training objective combines the KL loss with the standard language modeling loss:

$$\mathcal{L}_{Total} = \lambda \times \mathcal{L}_{LLM} + (1 - \lambda) \times \mathcal{L}_{KL} \quad (10)$$

where \mathcal{L}_{LLM} is the causal language modeling loss and λ balances the two objectives. To complete the alignment process, we perform a final inoculation step of our clinical predictor f_ϕ on a small sample of rewrites generated by our KL-trained EHR rewriter.

5 Experimental Design

5.1 Datasets

We use three EHR datasets in our study: MIMIC-IV [20], a publicly available dataset from Beth Israel Deaconess Medical Center.

MIMIC-IV contains clinical data from hospital and ICU stays, and is widely adopted in healthcare research. To assess the generalization of our approach, we additionally consider two other datasets: eICU [39], which gathers ICU data from many critical care units throughout the U.S. and EFEMERIS [16], a private clinical dataset containing the medical history of pregnant women as well as their neonatal outcomes regarding baby congenital malformation at birth.

5.2 Clinical tasks and evaluation metrics

Mortality prediction (MOR) predicts whether the patient will die in the next hospital visit, based on tuples from the current visit. Formally, $f : \{\cup_{t=1}^T \mathcal{T}_t^e(t)\} \mapsto y_t[\mathcal{T}_t^{e+1}]$ with $y_t[\mathcal{T}_t^{e+1}] \in \{0, 1\}$ denotes the patient’s mortality status.

Readmission prediction (RA) predicts if a patient will be readmitted into hospital within 15 days. Formally, $f : \{\cup_{t=1}^T \mathcal{T}_t^e(t)\} \mapsto y_t[\mathcal{T}_t^{e+1}]$ with $y_t[\mathcal{T}_t^{e+1}]$ is equal to 1 if $\Delta T(V^{e+1}, V^e) \leq 15$, 0 otherwise, with ΔT is a time-interval function.

Length-of-stay prediction (LOS) predicts whether the patient’s hospital stay will be longer than 7 days, using the first 48 hours of the current visit. Formally, $f : \{\cup_{t=1}^{\leq 48h} \mathcal{T}_t^e(t)\} \mapsto y_t[\mathcal{T}_t^e]$ with $y_t[\mathcal{T}_t^e] \in \{0, 1\}$ denotes if the stay is longer than 7 days.

Congenital malformation prediction (MALF) predicts whether a newborn will present major malformation based on maternal EHR data. Formally, $f : \{\mathcal{T}_j^1, \mathcal{T}_j^2, \mathcal{T}_j^3\} \mapsto y_j$, where \mathcal{T}_j^i represents the maternal EHR data from month j , and $y_j \in \{0, 1\}$ indicates the presence of major malformation at birth.

We prepare task-specific cohorts as follows. For MOR and RA tasks, we omit patients who are deceased in the current hospital admission. For the LOS prediction task, we exclude patients with a length-of-stay shorter than 48 hours. Table 1 shows the final cohort statistics per task. We followed previous work pre-processing [49] for data preparation across all datasets. First, we exclude patients with more than two ICU stays per hospital admission, and with negative ICU or hospital length-of-stay, and patients under 18 years old. We then build our EHR by selecting the following tables: (1) for MIMIC-IV, we use hosp/patients, hosp/admissions, hosp/diagnosis, hosp/labevents, hosp/microbiologyevents, hosp/prescriptions, and hosp/transfers; (2) for eICU, diagnosis, microlab, lab, patient, medication, and treatment; and (3) EFEMERIS, demographics and prescriptions tables. We omit tables capturing dense bedside monitor signals or with substantial overlap with other tables[20, 49]. Finally, we randomly split each dataset into train, val, test sets (80/10/10).

To evaluate prediction effectiveness, we use the standard metrics for binary classification problems: the Area Under the Receiver Operating Characteristic Curve (ROC) and the Area Under the Precision-Recall Curve (PRC) scores. We report mean scores and standard deviations computed via bootstrap sampling with replacement over 1000 iterations.

5.3 Baselines

We compare our approach with three main groups of baselines: (1) *EHR-oriented models*, including classical methods such as RNN [2], RETAIN [3], and GRASP [55], which require particular hand-crafted features pre-processing following established protocols [18, 49]. We also consider foundation models that can process raw EHR without particular pre-processing, including REMED [22], and Llemr

Table 1: Cohort statistics for each task and dataset.

Dataset	Task (s)	Size (m_s)	# Neg.	# Pos. (%)
MIMIC-IV	MOR	29091	28506	585 (2.01%)
	RA	29091	12738	16353 (52.21%)
	LOS	53771	32637	21134 (39.30%)
eICU	MOR	32750	31750	1000 (3.05%)
	RA	32750	31840	910 (2.78%)
	LOS	197174	185224	11950 (6.06%)
EFEMERIS	MALF	134843	131623	3220 (2.39%)

[49]; (2) *Serialized Classifiers*, which convert EHR data into text and fine-tune pre-trained language models with classification heads, including ClinicalBERT [14], Llama (Llama3-8B) [9], Qwen (Qwen 2.5-7b) [52], and ModernBERT (base) [48]; (3) *Rewrite then predict models*, which transform the original EHR input before classification. We evaluate all the proposed feature selection methods (π^h , π^d , π^r), including state-of-the-art ones π_{mt}^d [25], π_{mrm}^d [7], and π_{rfe}^d [11] as well as self-gen [46], an LLM-based feature selector using Qwen2.5-7B as the backbone LLM.

5.4 Implementation details

Hardware and software configurations. All training and evaluations are performed using CUDA 12.4, PyTorch 2.6.0 and the HuggingFace Transformers library. We train our models with 4 NVIDIA H100 80GB GPUs. To implement the baselines, we use PyHealth 1.1.6 framework [53] when available. For efficient rewrite generation, we use the vLLM 0.8.3 library for accelerated inference. We perform grid search optimization for learning rates across all experiments, exploring values in the range $\{1e-3, 1e-4, 1e-5, 1e-6\}$.

Feature selector (π). We implement the data-driven feature selection methods π^d using the scikit-learn feature selection library. For the heuristic-based approaches π^h , we rely on the available clinical normal range values provided within each dataset. Across all approaches, following [17], we select the top- $\kappa = 30\%$ of features based on their computed relevance scores, and we add missing features among the top 10, to complement EHR information.

EHR rewriter. We evaluate our framework using two recent instruction tuned LLMs: Llama3-8B and Qwen2.5-7B. We construct our $\mathcal{D}_{R,w}$ dataset by fixing a top- $k = 25\%$ for the quality threshold τ_s . We fine-tune the rewriter using LoRA [13] with learning rate $2e-4$, 3 training epochs, rank $r = 8$, LoRA $\alpha = 16$, dropout rate of 0.05, float16 precision, weight decay of 0.1, batch size of 4 and gradient accumulation steps of 4. For KL fine-tuning, we construct \mathcal{D}_{Pr}^s using $n_i = 8$ rewrites per patient and train up to 4000 steps using batch size 16, learning rate $2e-6$, $\kappa = 0.01$, $\tau \in \{0.01, 0.1, 0.2, 0.3\}$, and $\lambda \in \{0, 0.25, 0.5, 0.75\}$. Optimal parameters are selected based on task-specific evaluation every 1000 steps.

Clinical predictor. We employ identical hyperparameters for both the Scorer^{*} classifier and clinical predictor. We sample 20% of training patients to build \mathcal{D}_{sub}^s and generate \mathcal{D}_s^s with 3 additional rewrites per patient. Both models use ModernBERT-base as the encoder backbone with a classification head, learning rate $2e-5$,

Table 2: Results on the MIMIC-IV clinical tasks. Bold and Underline show best and 2nd best scores. Metrics multiplied by 100.

Model	MIMIC-IV					
	MOR		RA		LOS	
	ROC(↑)	PRC(↑)	ROC(↑)	PRC(↑)	ROC(↑)	PRC(↑)
<i>EHR-oriented models</i>						
(a) RNN (Choi et al., 2014 [2])	63.33 ± 3.4	3.38 ± 0.7	65.76 ± 1.2	72.49 ± 1.3	73.02 ± 0.7	58.01 ± 1.3
(b) RETAIN (Choi et al., 2016 [3])	59.13 ± 3.9	4.06 ± 1.8	64.57 ± 1.2	70.74 ± 1.4	72.91 ± 0.7	58.21 ± 1.3
(c) GRASP (Zhang et al., 2021 [55])	58.94 ± 3.6	2.68 ± 0.5	63.95 ± 1.2	68.92 ± 1.4	70.48 ± 0.8	53.85 ± 1.3
(d) Llmr (Wu et al., 2024 [49])	52.07 ± 4.1	3.17 ± 1.3	59.20 ± 1.1	64.89 ± 1.3	75.36 ± 0.7	62.21 ± 1.3
(e) REMed (Kim et al., 2024 [22])	52.40 ± 3.0	2.16 ± 0.5	68.26 ± 1.3	75.31 ± 1.9	80.15 ± 1.9	68.06 ± 2.9
<i>Serialized Classifiers</i>						
(f) BioMedBERT (Gu et al., 2020 [10])	64.24 ± 3.6	3.38 ± 0.7	70.36 ± 1.0	76.66 ± 1.0	77.59 ± 0.7	64.42 ± 1.3
(g) Qwen (Yang et al., 2024 [52])	53.87 ± 4.1	3.12 ± 0.9	61.06 ± 1.1	66.15 ± 1.3	73.78 ± 0.7	60.75 ± 1.3
(h) Llama (Dubey et al., 2024 [9])	59.45 ± 4.0	2.86 ± 0.6	64.93 ± 1.1	70.39 ± 1.3	75.09 ± 0.7	62.42 ± 1.2
(i) ModernBERT (Warner et al., 2024 [48])	58.01 ± 3.7	3.26 ± 1.2	70.98 ± 1.0	77.25 ± 1.0	<u>80.40 ± 0.6</u>	68.42 ± 1.2
<i>Rewrite-then-Predict models</i>						
(j) self-gen (Sui et al., 2024 [46])	58.38 ± 3.6	2.95 ± 0.9	69.19 ± 1.0	76.24 ± 1.0	79.02 ± 0.6	67.04 ± 1.2
(k) π_t^h	61.00 ± 4.2	3.58 ± 1.0	65.25 ± 1.0	73.77 ± 1.1	70.04 ± 0.7	54.74 ± 1.3
(l) π_v^h	56.05 ± 3.7	2.57 ± 0.5	64.47 ± 1.0	70.06 ± 1.3	66.63 ± 0.8	50.79 ± 1.3
(m) π_{ml}^d (Lewis et al., 1992 [25])	56.39 ± 3.8	2.93 ± 0.7	67.98 ± 1.0	74.57 ± 1.1	76.13 ± 0.7	60.57 ± 1.2
(n) π_{mmr}^d (Ding et al., 2005 [7])	60.00 ± 3.6	3.21 ± 0.8	63.08 ± 1.1	68.10 ± 1.3	79.48 ± 0.6	68.53 ± 1.2
(o) π_{rfe}^d (Guyon et al., 2002 [11])	62.54 ± 3.5	3.69 ± 1.3	62.03 ± 1.0	67.55 ± 1.3	76.84 ± 0.7	63.22 ± 1.2
(p) π_r^f	55.82 ± 3.7	2.63 ± 0.7	66.60 ± 1.0	72.34 ± 1.2	73.33 ± 0.7	59.29 ± 1.3
(q) π_v^f	59.37 ± 3.8	3.73 ± 1.5	68.03 ± 1.0	74.37 ± 1.1	77.51 ± 0.7	63.41 ± 1.2
ReToP _{Qwen} (ours)	<u>69.75 ± 3.8</u>	<u>4.61 ± 1.0</u>	<u>71.45 ± 0.9</u>	<u>77.92 ± 1.0</u>	80.38 ± 0.6	<u>68.91 ± 1.2</u>
ReToP _{Llama} (ours)	71.92 ± 3.6	6.20 ± 2.2	72.05 ± 0.9	78.25 ± 1.0	<u>80.62 ± 0.6</u>	<u>69.05 ± 1.2</u>

context length 8192 tokens, 10 training epochs with early stopping patience 3, batch size 32, and gradient accumulation steps 4. For the inoculation, we reduce the learning rate to $2e - 6$ and 4000 input samples to ensure optimal integration with the rewriter outputs.

6 Experimental Results and Analysis

6.1 Main results

Overall, Table 2 shows that the **ReToP** framework achieves a significant performance increase (t-test) over all the baselines and across all the clinical tasks. The performance increase ranges are respectively up to 38.12%, 33.51%, and 28.84% over respectively the *EHR-oriented*, *Serialized classifiers* and *Rewrite-Then-Predict* baseline models. Notably, we can see that among the *Rewrite-then-predict models*, all the feature selector operators ((k) . . . (q)) do not show consistent performance improvements. In contrast, our framework with the **ReToP_{Llama}** model achieves the highest performance across all the tasks, with improvements up to +23% compared to the best-performing baseline (ModernBERT). Specifically, **ReToP_{Llama}** obtains improvements of 13.92 for MOR, 1.07 RA, and 0.22 for LOS tasks. **ReToP_{Qwen}** shows similar improvements, demonstrating the robustness of our framework.

Finally, performance analysis across all the tasks raises an important observation: imbalanced tasks such as MOR with only 2% positive cases (§ Table 1), leverage greater benefit from the **ReToP** framework (+23%) compared to more balanced tasks like RA and LOS tasks with larger training sets (with nearly twice the size

of RA and MOR sets). This trend suggests that our rewrite-based approach provides greater value for rare clinical events, where high-quality synthetic representations can effectively address data scarcity, which is a common challenge in clinical prediction tasks.

6.2 Ablation study

We conduct comprehensive ablation studies to analyze the impact of different components of our **ReToP** framework at the training stage, with the following scenarios:

- (1) *w/o \mathcal{D}_{Rw}* : we replace our synthetic dataset generation \mathcal{D}_{Rw} with zero-shot LLM rewriting. Specifically, we replace \mathcal{D}_{Rw} by $\bigcup_{i \in \mathbb{S}} \{(\mathcal{P}_i, \mathcal{M}_i^g(\mathcal{P}_i | \mathcal{P}_i))\}_{1 \leq i \leq 4}$, $\mathcal{P}_i \in \mathbb{P}^S$, with 4 rewrites per EHR (w/o Algorithm 1).
- (2) *w/o Rewriter*: we use an off-the-shelf LLM, \mathcal{M}_i^g , instead of the \mathcal{M}_i EHR rewriter (w/o step 1, Fig. 2).
- (3) *w/o KL*: we train the LLM rewriter without the KL divergence loss, removing the alignment between the rewriter and the clinical prediction objectives (w/o step 3, Fig. 2).

We report our results in Table 3. Consistent with our main results, Llama3-8B demonstrates marginally superior performance compared to Qwen2.5-7B across all the ablation scenarios. We can observe that ablating the EHR alignment component (*w/o KL*) causes the most important performance degradation across both backbone LLMs, with the most pronounced effects on the imbalanced MOR prediction task. Specifically, we can see that substantial PRC drops of 31.5% and 41.8% for Qwen and Llama, respectively. Similarly, for the RA task, we observe a degradation between 4.7% and 6.3%.

Table 3: Ablation study results. Bold and Underline indicate the best and 2nd best performance respectively. Values in parentheses () indicate percentage degradation relative to the complete ReToP framework, if any. All metrics are multiplied by 100.

Scenario	Qwen2.5-7B						Llama3-8B					
	MOR		RA		LOS		MOR		RA		LOS	
	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC
w/o D_{R_w}	64.92(6.9%)	5.22	69.49(2.7%)	76.11(2.3%)	78.78(2.0%)	66.42(3.6%)	70.03(2.6%)	4.48(27.7%)	68.88(4.4%)	75.38(3.7%)	76.98(4.5%)	65.08(5.7%)
w/o Rewriter	58.83(15.7%)	3.16(31.5%)	68.69(3.9%)	74.67(4.2%)	79.06(1.6%)	67.18(2.5%)	65.42(9.0%)	4.93(20.5%)	68.86(4.4%)	73.48(6.1%)	79.37(1.6%)	67.27(2.6%)
w/o KL	<u>64.93(6.9%)</u>	3.16(31.5%)	67.22(5.9%)	74.25(4.7%)	78.29(2.6%)	66.44(3.6%)	59.00(18.0%)	3.61(41.8%)	67.40(6.5%)	73.33(6.3%)	79.09(1.9%)	<u>67.38(2.4%)</u>
ReToP	69.75	<u>4.61</u>	71.45	77.92	80.38	68.91	71.92	6.20	72.05	78.25	80.62	69.05

Table 4: Transferability performance with ROC/PRC metrics.

Model	MIMIC-IV \rightarrow eICU			\rightarrow EFEMERIS	
	MOR	RA	LOS	MALF	
ModernBERT	81.30/21.86	86.71/23.80	93.19/56.67	49.79/6.44	
ReToP _{Llama}	81.42/22.29	86.72/23.78	93.16/56.78	51.53/6.93	
ReToP _{Qwen}	81.40/21.93	86.79/23.91	93.17/56.80	50.41/7.08	

For the LOS task, we observe a more modest degradation between 2.4% and 3.6%, suggesting that KL training provides critical value, particularly for tasks with severe class imbalance.

Ablating the fine-tuning of the EHR rewriter (*w/o Rewriter*) follows similar degradation patterns, with the MOR task suffering from a decrease of 31.5% and 20.5%, for Qwen and Llama in PRC scores, respectively. This proves that fine-tuning the EHR rewriter is an equally crucial step in our pipeline.

Finally, we observe that removing our synthetic training dataset (*w/o D_{R_w}*) has the most pronounced impact on the LOS task, indicating that pseudo-label augmentation provides greater value for tasks with larger training sets where diverse synthetic examples can capture broader clinical patterns.

6.3 Model transferability

We evaluate the transferability of the ReToP framework on unseen data, namely, eICU¹ and unseen task, namely, MALF with the EFEMERIS dataset. We use the rewriter of each corresponding task². We adopt a low-cost adaptation approach, implementing efficiency measures at both model and data levels. Rather than fine-tuning the complete ReToP framework ($> 8B$ params.), we train only the clinical predictor with 149M params. (1.8% of total). For data efficiency, we limit rewritten data to the test set and 4000 training samples (equivalent to 3.6% of the largest dataset) for predictor inoculation. We use our best baseline classifier (ModernBERT). Due to computational constraints, we evaluate on 20% of the test set for LOS and MALF tasks. As shown in Table 4, ReToP_{Qwen} and ReToP_{Llama} improve PRC scores up to +1.97%, +0.46%, and +9.94% for MOR, RA, and MALF tasks, correspondingly, while maintaining competitive LOS performance. These results demonstrate generalization with minimal adaptation cost, particularly for highly imbalanced tasks (MOR, RA, and MALF) with up to 3.05% positive examples.

¹We defined our clinical task at the ICU visit instead of the hospital visit following [53]

²For MALF, we select the best EHR rewriter among the 3 tasks.




Figure 3: ROC scores across α values for inference level (§ Eq. 2) using ReToP_{Llama} on MIMIC-IV, shown for MOR (left), RA (center), and LOS (right) tasks. Curves are stratified by EHR length: overall (blue), short EHR < 2048 (green), medium EHR in [2048, 4096] (yellow), and long > 4096 (red).

6.4 Model analysis

Effect of interpolated prediction. We evaluate the impact of the interpolation parameter α (§ Eq. 2) for leveraging original EHRs and rewrites at the inference stage. Figure 3 shows ROC performance across different α values (blue curves), where $\alpha = 0$ represents the inference with no rewrites, and $\alpha = 1$, using the EHR rewriter. Aligned with our previous findings, optimal values are task dependent. MOR benefits more from rewrite-only input, showing that our EHR rewriter mitigates sparsity and noise characteristics of mortality-related EHR data. Conversely, RA and LOS benefit from no rewrites ($\alpha = 0$), suggesting that preserving full clinical detail is more valuable than noise reduction. Figure 3 further stratifies performance by EHR length. For the MOR task, rewrites ($\alpha = 1$) consistently outperforms across all the input lengths. However, for RA and LOS, short EHRs show little differences between original and rewritten versions, while medium and longer EHRs exhibit degraded performance when using rewrites, suggesting that rewrite quality decreases with increasing EHR length for these tasks.

Effect of KL training. We analyze KL divergence training across training steps and balancing parameter λ (§ Eq. 10) using ReToP_{Llama} without inoculation. Figure 4 shows that that models with $\lambda \geq 0.5$ do not further improve after 3000 – 4000 steps, suggesting early stopping. For $\lambda \leq 0.25$, training is noisy with early overfitting except for RA. MOR is more sensitive to λ , benefiting from stronger regularization early on, while RA and LOS remain robust across λ values. Our results suggest regularization should consider class imbalance.




Figure 4: Effect of λ on KL training (up to 4k steps, eval. every 1000 steps) for MOR, RA, LOS using ReToP_{Llama}.

7 Case Study

Our aim here is to check and then get insights into the faithfulness and clinical value under the expert perspective of the best ReToP setting. To achieve this goal, we qualitatively analyze a sample of 15 patient EHRs where ReToP accurately predicts the MOR task. For each case, we collect rewrites from self-gen, ReToP w/o KL, and full ReToP using Llama backbone model. Annotators blindly evaluate each rewrite on two criteria using three labels ('Yes', 'Partially', 'No'): (1) Faithfulness: whether the rewrite is entailed by the original EHR, judged by three reviewers including one expert, and (2) Actionability: whether the rewrite includes clinical features that likely support the clinical decision-making, annotated by one expert.

Table 5 shows the results in terms of ratios for each label and each model on Faithfulness and Actionability criteria. Overall, ReToP w/o KL shows higher faithfulness than ReToP and self-gen (based on the 'Y' annotation) and lower faithfulness after KL, leading to a decrease of the 'Y' and an increase of the 'P' and 'N' annotations. This trend is consistent with the expert annotation on the actionability, showing an opposite pattern on Actionability criteria. We can see that self-gen is more likely to include expected predictive features by the experts than the w/o KL model and, more importantly, the full ReToP scenario, with nearly half of the annotations (46) revealing the increasing presence of unexpected features for the expert. All these results bring two important insights of the ReToP framework: (1) the EHR rewriter (w/o KL) intrinsically outputs EHR rewrites that are faithful to the original ones, as targeted in addressing challenge C1; (2) the KL alignment between the EHR rewriter and clinical predictors revises the rewriter by emphasizing predictive clinical features, in addressing C2, but which are seemingly serendipitous for the expert regarding the clinical task at hand.

We dig into these results with a qualitative analysis on Figure 5. Examining faithfulness across all the rewrites, self-gen preserves all the original features but introduces unfaithful information (highlighted in yellow). For instance, while the original EHR indicates "Alcohol abuse, unspecified, the self-gen rewrite adds complementary information stating that "... alcohol abuse may impact his medication, not present in the source data. In contrast, both ReToP-based rewrites lean to filter features rather than adding new spans, which preserves their faithfulness. Interestingly, we can see that ReToP w/o KL reduces diagnoses from 30 to 24 and laboratory tests from 51 to 40, while ReToP further reduces tests to 34, retaining only task-relevant features that optimize prediction. This would explain the decrease in actionability according to the expert perspective and opens avenues of research for model explainability about the underlying reasons to filter.

Table 5: Qualitative analysis of ReToP (w/o and w/ KL models) vs. self-gen on MOR based on faithfulness and actionability.

Model	% Faithfulness			% Actionability		
	Y	P	N	Y	P	N
self-gen	60	40	0	80	20	0
ReToP (w/o KL)	73	27	0	0	100	0
ReToP	71	22	7	7	47	46



Figure 5: Example patient EHR for the MOR task. Red text shows main content-based differences across rewrites, while highlighted text shows generated unfaithful information.

8 Conclusion

We introduced ReToP, a new framework that leverages LLMs to enhance clinical prediction performance. ReToP trains an LLM-based EHR rewriter using synthetic EHR rewrites built upon health-related feature selection methods. Then, ReToP refines the EHR rewriter through an end-to-end training guided by the clinical predictor supervision using a KL loss. ReToP significantly outperforms a set of state-of-the-art baselines across representative clinical prediction tasks. Our proposed framework exhibits reasonable transfer ability to out-of-domain datasets and tasks. By designing EHR rewriters that can be efficiently aligned with downstream clinical tasks, ReToP opens up potential directions for effective healthcare AI systems. Future work could explore extending this framework to a wider range of clinical tasks, including multi-label classification tasks, and investigating the right compromise between model performance and model explainability for domain experts.

Acknowledgments

This work has been supported by the In-Utero project funded by HDH (France) and FRQS (Canada). This work was also granted access to the HPC resources of IDRIS under the allocation 2025-AD011015371R1 made by GENCI.

Ethical Considerations

In this work, we used available de-identified datasets from the medical domain, including MIMIC-IV, eICU, and EFEMERIS datasets, with proper attribution to sources. These datasets contain patient data that has been anonymized and de-identified by established privacy protection standards, ensuring no individual patient information can be traced or identified.

While our ReTOP framework demonstrates improved clinical prediction performance, we emphasize that these predictions should be exclusively used as decision-supporting tools for experts.

References

- [1] Jawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024. Benchmarking Large Language Models in Retrieval-Augmented Generation. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 16 (Mar. 2024), 17754–17762. <https://doi.org/10.1609/aaai.v38i16.29728>
- [2] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the Properties of Neural Machine Translation: Encoder–Decoder Approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Dekai Wu, Marine Carpuat, Xavier Carreras, and Eva Maria Vecchi (Eds.). Association for Computational Linguistics, Doha, Qatar, 103–111. <https://doi.org/10.13151/v1/W14-4012>
- [3] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. 2016. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. *Advances in neural information processing systems* 29 (2016), 3512–3520.
- [4] Edward Choi, Cao Xiao, Walter F. Stewart, and Jimeng Sun. 2018. MiME: multi-level medical embedding of electronic health records for predictive healthcare. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (Montreal, Canada) (NIPS'18)*. Curran Associates Inc., Red Hook, NY, USA, 4552–4562.
- [5] Edward Choi, Zhen Xu, Yujia Li, Michael Dusenberry, Gerardo Flores, Emily Xue, and Andrew Dai. 2020. Learning the Graphical Structure of Electronic Health Records with Graph Convolutional Transformer. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 01 (Apr. 2020), 606–613. <https://doi.org/10.1609/aaai.v34i01.5400>
- [6] Kristy Choi, Chris Cundy, Sanjari Srivastava, and Stefano Ermon. 2022. LMPriors: Pre-Trained Language Models as Task-Specific Priors. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- [7] Chris Ding and Hanchuan Peng. 2005. Minimum redundancy feature selection from microarray gene expression data. *Journal of bioinformatics and computational biology* 3, 02 (2005), 185–205.
- [8] Dmitry Dligach and Timothy Miller. 2018. Learning Patient Representations from Text. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*: 119–123.
- [9] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv e-prints* (2024), arXiv:2407.
- [10] Yu Gu, Robert Timn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. *ACM Trans. Comput. Healthcare* 3, 1, Article 2 (Oct. 2021), 23 pages. <https://doi.org/10.1145/3458754>
- [11] Isabelle Guyon, Jason Weston, Stephen Barnhill, and Vladimir Vapnik. 2002. Gene selection for cancer classification using support vector machines. *Machine learning* 46, 1 (2002), 389–422.
- [12] Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David A. Sontag. 2022. TabLLM: Few-shot Classification of Tabular Data with Large Language Models. In *AISTAT*, Vol. abs/2210.10723.
- [13] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR* 1, 2 (2022), 3.
- [14] Kexin Huang, Jaan Allosaar, and Rajesh Ranganath. 2019. ClinicalBERT: Modeling Clinical Notes and Predicting Hospital Readmission. *arXiv:1904.05342* (2019).
- [15] Kyungho Hur, Jungwoo Oh, Junu Kim, Jiyoum Kim, Min Jae Lee, Eunbyeo Ch, Seong-Eun Moon, Young-Hak Kim, Louis Atallah, and Edward Choi. 2023. GenHP: General healthcare predictive framework for multi-task multi-source learning. *IEEE Journal of Biomedical and Health Informatics* 28, 1 (2023), 502–513.
- [16] Lacroix I, Hurauli C, Sarramon MF, Guillard C, Berrebi A, Grau M, Abouy-Cossard C, Bourrel R, Elefant E, Monsttraue JL, and Damase-Michel C. 2009. Prescription of drugs during pregnancy: a study using EFEMERIS, the new French database. *European journal of clinical pharmacology* 65, 8 (2009), 839–846.
- [17] Daniel P Jeong, Zachary Chase Lipton, and Pradeep Kumar Ravikumar. 2025. LLM-Select: Feature Selection with Large Language Models. *Transactions on Machine Learning Research* (2025).
- [18] Pengcheng Jiang, Cao Xiao, Adam Richard Cross, and Jimeng Sun. 2024. GraphCare: Enhancing Healthcare Predictions with Personalized Knowledge Graphs. In *The Twelfth International Conference on Learning Representations*.
- [19] Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. How Can We Know Language Models Know? On the Calibration of Language Models for Question Answering. *Transactions of the Association for Computational Linguistics* 9 (2021), 962–977. <https://doi.org/10.1162/tacl.2021.04007>
- [20] Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Geyers, Ayad Shammout, Steven Hong, Tom J Pollard, Sicheong Hao, Benjamin Moody, Brian Gow, et al. 2023. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data* 10, 1 (2023), 1.
- [21] Stone K, Zwiggelaar R, Jones P, and Mac Parthalain N. 2022. A systematic review of the prediction of hospital length of stay: Towards a unified framework. *PLOS Digit Health* 14 (2022). <https://doi.org/10.1371/journal.pdig.0000017>
- [22] Junu Kim, Chaeun Shim, Bosco Seong Kyu Yang, Chami Im, Sung Yoon Lim, Han-Gil Jeong, and Edward Choi. 2024. General-Purpose Retrieval-Enhanced Medical Prediction Model Using Near-Infinite History. In *Proceedings of the 9th Machine Learning for Healthcare Conference (Proceedings of Machine Learning Research, Vol. 252)*. Kaivalya Deshpande, Madalina Fiterau, Shalmali Joshi, Zachary Lipton, Rajesh Ranganath, and Inigo Urteaga (Eds.), PMLR.
- [23] Sunjun Kwon, Junu Kim, Jiyoum Kim, Sujong Im, Eunbyeo Cho, Seongsu Bae, Jungwoo Oh, Gyubok Lee, Jong Hak Moon, Seung Chan You, Seungjin Bae, Chang Hoon Han, Yoon Bin Jung, Yohan Jo, and Edward Choi. 2024. Publicly Shareable Clinical Large Language Model Built on Synthetic Clinical Notes. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 5148–5168. <https://doi.org/10.18653/v1/2024.findings-acl305>
- [24] I. Landi, B.S. Glicksberg, and H.C. et al. Lee. 2020. Deep representation learning of electronic health records to unlock patient stratification at scale. *npi Digit. Med.* 3, 3 (Jul. 2020). <https://ojs.aaai.org/index.php/AAAI/article/view/29728>
- [25] David D Lewis. 1992. Feature selection and feature extraction for text categorization. In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*.
- [26] Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, and Michael Bendersky. 2024. Learning to Rewrite Prompts for Personalized Text Generation. In *Proceedings of the ACM Web Conference 2024 (WWW '24)*. ACM, 3367–3378. <https://doi.org/10.1145/3589334.3645408>
- [27] Dawei Li, Zhen Tan, and Huan Liu. 2025. Exploring large language models for feature selection: A data-centric perspective. *ACM SIGKDD Explorations Newsletter* 26, 2 (2025), 44–53.
- [28] Rumeng Li, Xun Wang, and Hong Yu. 2024. LlamaCare: An Instruction Fine-Tuned Large Language Model for Clinical NLP. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakiti, and Chianwen Xue (Eds.), ELRA and ICCL, Torino, Italia, 10632–10641. <https://aclanthology.org/2024.lrec-main.930/>
- [29] Y. Li, S. Rao, J. Solares, A. Hassaine, R. Ramakrishnan, D. Canoy, Y. Zhu, K. Rahimi, and G. Salimi-Khorshidi. 2020. BEHRT: Transformer for Electronic Health Records. , 7155 pages. <https://doi.org/10.1038/s41598-020-62922-y>
- [30] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual Instruction Tuning.
- [31] Haochen Liu, Song Wang, Yaochen Zhu, Yushun Dong, and Jundong Li. 2024. Knowledge Graph-Enhanced Large Language Models via Path Selection. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 6311–6321. <https://doi.org/10.18653/v1/2024.findings-acl376>
- [32] Qidong Liu, Xian Wu, Xiangyu Zhao, Yuanshao Zhu, Derong Xu, Feng Tian, and Yefeng Zheng. 2024. When MOE Meets LLMs: Parameter Efficient Fine-tuning for Multi-task Medical Applications. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (Washington DC, USA) (SIGIR '24)*. Association for Computing Machinery, New York, NY, USA, 1104–1114. <https://doi.org/10.1145/3626772.3657722>
- [33] Jesús Lovón-Melgarejo, Martín Mouysset, Jo Olievan, José G. Moreno, Christine Damase-Michel, and Lynda Tamme. 2025. Evaluating LLM Abilities to Understand Patient Electronic Health Records: A Comprehensive Study of Patient Data Extraction and Retrieval. In *European Conference in Information Retrieval (ECIR)*. Elsevier.
- [34] Sushil Madhumita, ester Simon, Lukyckx Kim, and Daelemans Walter. 2018. Patient representation learning and interpretable evaluation using clinical notes. *Journal of Biomedical Informatics* 84 (2018), 103–113.
- [35] Fengran Mo, Jian-Yun Nie, Kaiyu Huang, Kelong Mao, Yutao Zhu, Peng Li, and Yang Liu. 2023. Learning to Relate to Previous Turns in Conversational Search. In *29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (SIGKDD)*.
- [36] Aakanksha Nair, Sravanthi Parasa, Sergey Feldman, Lucy Lu Wang, and Tom Hope. 2022. Latent-Augmented Clinical Outcome Prediction. In *Findings*

- of the Association for Computational Linguistics: NAACL 2022, Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, Seattle, United States, 438–453.
- [37] Tuan Dung Nguyen, Thanh Trung Huynh, Minh Hieu Phan, Quoc Viet Hung Nguyen, and Phi Le Nguyen. 2024. CARER – ClinicAI Reasoning-Enhanced Representation for Temporal Health Risk Prediction. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 10392–10407. <https://doi.org/10.18653/v1/2024.emnlp-main.580>
- [38] 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)*. Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 2463–2473. <https://doi.org/10.18653/v1/D19-1250>
- [39] Tom J Pollard, Alistair EW Johnson, Jesse D Raffa, Leo A Celi, Roger G Mark, and Omar Badawi. 2018. The eICU Collaborative Research Database, a freely available multi-center database for critical care research. *Scientific data* 5, 1 (2018), 1–13.
- [40] Laila Rasmay, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. 2020. MedBERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ Digital Medicine* 4 (2020). <https://api.semanticscholar.org/CorpusID:21889776>
- [41] Devendra Singh Sachan, Mike Lewis, Dan Yogatama, Luke Zettlemoyer, Joelle Pineau, and Manzil Zaheer. 2023. Questions Are All You Need to Train a Dense Passage Retriever. *Transactions of the Association for Computational Linguistics* 11 (2023), 600–616. https://doi.org/10.1162/tacl_a_00564
- [42] Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suh. 2024. Quantifying Language Models’ Sensitivity to Spurious Features in Prompt Design: or How I learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=RhslyNXJT>
- [43] Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. REPLUG: Retrieval-Augmented Black-Box Language Models. In *NAACL-HLT*.
- [44] Wenqi Shi, Ran Xu, Yuchen Zhuang, Yue Yu, Jiayu Zhang, Hang Wu, Yuanda Zhu, Joyce C. Ho, Carl Yang, and May Dongmei Wang. 2024. EHRAGENT: Code Empowers Large Language Models for Few-shot Complex Tabular Reasoning on Electronic Health Records. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 22315–22339. <https://doi.org/10.18653/v1/2024.emnlp-main.1245>
- [45] Ananya Singha, José Cambronero, Sumit Gulwani, Vu Le, and Chris Parnin. 2023. Tabular Representation, Noisy Operators, and Impacts on Table Structure Understanding Tasks in LLMs. In *Table Representation Learning Workshop at NeurIPS 2023*.
- [46] Yuan Sui, Mengyu Zhou, Mingjie Zhou, Shi Han, and Dongmei Zhang. 2024. Table Meets LLM: Can Large Language Models Understand Structured Table Data? A Benchmark and Empirical Study. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (, Merida, Mexico.) (WSDM '24)*. Association for Computing Machinery, New York, NY, USA, 645–654. <https://doi.org/10.1145/3616855.3635752>
- [47] Lynda Tamine and Lorraine Goeuriot. 2021. Semantic Information Retrieval on Medical Texts: Research Challenges, Survey, and Open Issues. *ACM Comput. Surv.* 54, 7, Article 146 (Sept. 2021), 38 pages. <https://doi.org/10.1145/3462476>
- [48] Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Welal, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. 2024. Smarter, Better, Faster, Longer: A Modern Bidirectional Encoder for Fast, Memory Efficient, and Long Context Finetuning and Inference. arXiv:2412.13663 [cs.LG]. <https://arxiv.org/abs/2412.13663>
- [49] Zhenbang Wu, Anant Dadu, Mike Nalls, Faraz Faghri, and Jimeng Sun. 2024. Instruction Tuning Large Language Models to Understand Electronic Health Records. In *Advances in Neural Information Processing Systems*. A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 54772–54786.
- [50] Cao Xiao, Edward Choi, and Jimeng Sun. 2018. Opportunities and challenges in developing deep learning models using electronic health records data: a systematic review. *JAMIA* 25, 10 (2018), 1419–1428.
- [51] Ran Xu, Wenqi Shi, Yue Yu, Yuchen Zhuang, Bowen Jin, May Dongmei Wang, Joyce Ho, and Carl Yang. 2024. RAM-EHR: Retrieval Augmentation Meets Clinical Predictions on Electronic Health Records. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 754–765.
- [52] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jilin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Li, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuyong Liu, Zeyu Chu, Zhenru Zhang, and Zhihao Fan. 2024. Qwen2 Technical Report. arXiv preprint arXiv:2407.10671 (2024).
- [53] Chaoyi Yang, Zhenbang Wu, Patrick Jiang, Zhen Lin, Junyi Gao, Benjamin Daneek, and Jimeng Sun. 2023. PyHealth: A Deep Learning Toolkit for Healthcare Predictive Modeling. In *Proceedings of the 27th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) 2023*. <https://github.com/sunlabusc/PyHealth>
- [54] Zhichao Yang, Avijit Mitra, Weisong Liu, Dan Berlowitz, and Hong Yu. 2023. TransformEHR: transformer-based encoder-decoder generative model to enhance prediction of disease outcomes using electronic health records. *Nature Communications* 14 (2023). <https://api.semanticscholar.org/CorpusID:265503777>
- [55] Zhaohao Zhang, Xin Gao, Liantao Ma, Yasha Wang, Jiangtao Wang, and Wen Tang. 2021. GRASP: generic framework for health status representation learning based on incorporating knowledge from similar patients. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35, 715–723.
- [56] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate Before Use: Improving Few-shot Performance of Language Models. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.), PMLR, 12697–12706.
- [57] Yinghao Zhu, Changyu Ren, Zixiang Wang, Xiaochen Zheng, Shiyun Xie, Junlan Feng, Xi Zhu, Zhoujun Li, Liantao Ma, and Chengwei Pan. 2024. EMERGE: Enhancing Multimodal Electronic Health Records Predictive Modeling with Retrieval-Augmented Generation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (Boise, ID, USA) (CIKM '24)*. Association for Computing Machinery, New York, NY, USA, 3549–3559. <https://doi.org/10.1145/3627673.3679582>
- [58] Yinghao Zhu, Zixiang Wang, Long He, Shiyun Xie, Xiaochen Zheng, Liantao Ma, and Chengwei Pan. 2024. PRISM: Mitigating EHR Data Sparsity via Learning from Missing Feature Calibrated Prototype Patient Representations (CIKM '24). Association for Computing Machinery, New York, NY, USA, 3560–3569. <https://doi.org/10.1145/3627673.3679521>