PANCDetect: Early Detection of Pancreatic Cancer from Multi-modal EHR data with LLM Embeddings

Zicheng Jin[1*], Xuhui Guo[1*], Zehua Wang[1,2*], Qiang Yang[3], Xiaotong Yang[1], Xinyu Zhang[4], Rui Yin[3], Lana X. Garmire[1$]

[1]Department of Computational Medicine and Bioinformatics, University of Michigan, Ann Arbor, MI, USA

[2]Current Address: Department of Computer Science, University of Virginia, Charlottesville, VA, USA

[3]Department of Health Outcomes & Biomedical Informatics, University of Florida, Gainesville, FL, USA

[4]Department of Mathematics, University of Michigan, Ann Arbor, MI, USA

[*] These authors contributed equally to this work

[$] Corresponding author email: lgarmire@umich.edu

Abstract

Background

Pancreatic cancer (PANC) is often diagnosed at late stages due to the absence of specific early symptoms, resulting in one of the highest cancer mortality rates. While imaging modalities such as MRI and CT offer high diagnostic accuracy, their population-wide application is however impractical due to the cost. Electronic health records (EHRs) provide a routine, easily accessible, longitudinal and scalable data source for risk prediction, particularly for diseases with no specific symptom such as PANC.

Method

We introduce PANCDetect, a multimodal framework that leverages large language model (LLM)-derived embeddings of diagnoses, procedures, medications, and laboratory tests, and integrates these data modalities through a Transformer-based architecture. We train the model on MarketScan (≈250M patients), and validate it externally on additional large real-world EHR datasets of University of Michigan Precision Health, or UMPH data (n≈6M patients) and OneFlorida+ data (n≈26M patients). We then fine-tuned the general model on UMPH EHR data. We evaluated the performance of both models using metrics including area under the receiver operating characteristic curve (AUROC) and area under the precision-recall-gain curve (AUPRG). We assessed the top predictive features with integrated gradients (IG).

Result

In the MarketScan cohort, PANCDetect achieved an AUROC of 0.812 and AUPRG of 0.851 at the 6-month prediction window, and an AUROC of 0.735 and AUPRG of 0.629 for 60-month prediction, significantly outperforming CancerRiskNet. External validation on UMPH and OneFlorida+ demonstrated good generalizability, with 6-months AUROC scores of 0.7111 and 0.793, respectively. Fine-tuning on UMPH with laboratory data further improved performance,

reaching an AUROC of 0.927 and an AUPRG of 0.979 at 6 months. Even at the 60-month horizon, the refined PANCDetect model maintained strong performance, with an AUROC of 0.927 and AUPRG of 0.979. Attribution analysis highlighted type 2 diabetes, pancreatic diseases, personal and family cancer history as the most important risk factors.

Conclusion

PANCDetect is the state-of-the-art method integrating multimodal EHR data with LLM embeddings for accurate, interpretable, and generalizable early prediction of pancreatic cancer. This framework holds promise for precision screening of high-risk patients, with the potential to improve survival outcomes without increasing healthcare costs.

Introduction

Pancreatic cancer (PANC), is characterized by the lack of specific early clinical manifestations.[1] In many cases, individuals may experience vague or nonspecific symptoms that precede diagnosis by years.[2] Based on a global analysis of data from 48 countries, incidence and mortality of PANC are rising globally, particularly among women and individuals aged 50 and older, with concerning increases also seen in younger populations.[3] The absence of validated screening protocols, lack of reliable early biomarkers, combined with the retroperitoneal and anatomically inaccessible location of the pancreas, hinders timely detection. As a result, only a small proportion of patients are diagnosed at an early, resectable stage, resulting in persistently high mortality rate of this malignancy.[4] Additionally, delayed diagnosis of PANC imposes substantial financial burdens on patients, families, and healthcare systems, as late-stage treatments are far more costly than early surgery with adjuvant therapy.[5]

Routine MRI and CT as an effective screening method can detect PANC with high sensitivity and specificity,[6] but applying cross-sectional imaging at population scale is impractical due to prohibitive costs and resource burden.[7] Electronic health record (EHR) is a much more economical, accessible and scalable data resource. EHR data include valuable baseline information such as demographics, as well as the longitudinal information including disease diagnoses, laboratory results, medication, and procedures, make them well suited for early-detection and prognosis modeling.[8] Moreover, the standardized EHR frameworks allows them to be more AI ready, and comparable across different healthcare systems, which enhances the generalizability of predictive models across health centers.[9–11]

Early transformer-based pretrained models demonstrated the feasibility of applying self-attention architectures to structured EHR data for general prediction of future diagnoses of patients. Models such as BEHRT[12], Hi-BEHRT[13], and Med-BERT[14] adapted BERT-style encoders to longitudinal

claims and clinical coding sequences, effectively capturing temporal dynamics, disease progression, and patient heterogeneity. Building on this trend, researchers have begun applying Transformers to pancreatic cancer risk prediction, leveraging self-attention mechanisms to better model temporal dependencies and disease evolution. Moreover, recent advances in large language models (LLMs) offer new opportunities to extract clinically meaningful patterns hidden among the millions of patients' EHR data. By treating EHRs as sequences of "clinical tokens", such as diagnoses (ICD/SNOMED), procedures (CPT), medications (RxNorm/NDC), and discretized lab values, pretrained LLMs can learn rich contextualized embeddings that capture clinical semantics and longitudinal continuity across care settings.[15,16] General-purpose LLMs such as GPT have been adapted to biomedical corpora (e.g., BioGPT[17]) and even to institution-scale EHR data (e.g., NYUTron[18]), enabling predictive modeling, free-text understanding, and zero or few-shot adaptation to diverse downstream tasks.

A few recent studies have leveraged EHR data to build models tailored for predicting PANC risk, enabling identification of high-risk patient cohort[19–22]. Among these efforts, CancerRiskNet[21] stands out as a representative study that demonstrated the feasibility of using Transformers in this domain. However, its multi-window prediction design, which does not exclude disease events occurring shortly before cancer diagnosis, carries the risk of data leakage, potentially inflating accuracy estimates. Moreover, CancerRiskNet relies solely on diagnosis codes, and the authors acknowledged that incorporating other clinically informative modalities could further improve prediction accuracy with the real-world availability of data. Many existing models are trained on relatively small datasets, restricting generalizability across diverse populations and health systems[22]. To address these limitations and take advantage of LLMs, we propose PANCDetect, a Transformer-based multimodal prediction framework trained on the MarketScan dataset, which contains over 110,000 pancreatic cancer patients, making it the largest resource of its kind to date (Fig. 1). PANCDetect not only incorporates diagnostic information but also procedures,

medications, and laboratory tests into LLM-based embeddings. We also employed a more rigorous exclusion strategy to minimize data leakage in modeling and improve reliability. To validate the generalizability of our approach, we tested and fine-tuned the model across two independent healthcare systems. Consequently, PANCDetect substantially enhances prediction accuracy and clinical applicability for early pancreatic cancer detection as compared to other earlier efforts.

Methods

Study Cohorts

This study is conducted on three different retrospective datasets: Merative MarketScan Databases,[23] University of Michigan Precision Health (UMPH),[24] and OneFlorida+ Clinical Research Network.[25] MarketScan is a commercial database, which we have access to through the University of Michigan. UMPH and OneFlorida+ data usage was approved by the Institutional Review Board (IRB) of the University of Michigan Medical School (HUM#00259964) and the IRB of the University of Florida (#IRB202400752), with a waiver of informed consent for secondary use of de-identified EHR data. No data from UMPH or OneFlorida+ was shared outside the respective organizations, ensuring compliance with data security and confidentiality.

The MarketScan Database contains records from approximately 250 million patients from 2009–2022. We identified 113,518 individuals diagnosed with PANC (see Supplementary Table 1 for the list of diagnosis codes), with 50.8% males and 49.2% females (Table 1). The highest proportion of diagnoses is observed between ages 55 and 65 (Fig. 2A). As shown in Fig. 2B, most PANC patients exhibit trajectory lengths of 0–6 years with 50–400 accumulated diagnosis records, reflecting intensive healthcare utilization and multiple comorbidities in the years preceding diagnosis. In contrast, non-PANC patients generally have shorter trajectories.

The UMPH includes approximately six million patients from 2006 to 2025, among whom 8,257 individuals are recorded as PANC, including 53.3% male and 46.6% female (Table 1). The highest proportion of diagnoses is observed in ages of 65–70 (Fig. 2C). Compared to the MarketScan database, PANC patients in UMPH generally have shorter trajectory lengths with most of them spanning 2 years or less, reflecting the high volume of referred high-risk patients in a research hospital (Fig. 2D). Within this shorter time span, patients accumulated a large number of diagnosis records, reflecting higher encounter density and more clinical details.

OneFlorida+ Clinical Research Network encompasses around 26 million patients from 72 hospitals across Florida, Alabama, Georgia, and Arkansas. We identified 10,500 individuals diagnosed with PANC during 2013–2024, with 52.7% male and 47.3% female (Table 1). The majority of diagnoses were concentrated in the 65–75 age group (Fig. 2E), similar to the UMPH dataset. OneFlorida+ cohort showed that PANC patients had trajectory lengths of ≤4 years (Fig. 2F) and accumulated a median of 96 diagnosis records (Table 1). By contrast, non-PANC patients generally had fewer diagnosis records, consistent with lower healthcare utilization.

After identifying pancreatic cancer patients across the three datasets, we constructed the final case–control cohorts by applying a 1:5 matching strategy based on sex, age, and geographic location (Supplementary Fig. 1) to mitigate the substantial class imbalance. We then used the MarketScan cohort as the primary training set given its largest PANC patient size. We performed external validation on UMPH and OneFlorida+, and fine-tuned the model on the UMPH cohort for better adaptation to the local population characteristics.

Feature Embedding

As shown in Fig. 3A, the input features comprise time-varying information and time-invariant features. Time-varying information includes EHR clinical events such as diagnoses, procedures, and medications, each annotated with temporal markers such as visit date and patient age at that

visit. Time-invariant attributes include sex and geographic location. For the time-variant features, standardized textual descriptions of medical codes are encoded using the OpenAI text-embedding-3-small model to generate dense semantic representations (without patient private information). These embeddings are combined with positional embeddings to capture temporal order and irregular intervals after linear projection, thereby enhancing the representation of longitudinal patient trajectories. During fine-tuning on the UMPH cohort, we integrated laboratory test results as a new input modality, using the same embedding generation method to maintain consistency while enabling exploration of richer multi-modal inputs in a real-world clinical setting. Other time-invariant features (e.g., sex and geographic location) are mapped into dense representations using a feed-forward layer before being concatenated with the longitudinal sequence representations.

Multi-modal Fusion Strategy

The proposed framework integrates multi-modal information within a Transformer-based architecture (Fig. 3A). Here, three modalities, namely diagnoses, procedures, and medications, are modeled through self-attention blocks to capture intra-modality dependencies. Cross-attention mechanisms are then applied, using diagnosis events as queries and procedure and medication events as key–value inputs, enabling the model to learn complementary information across modalities. During fine-tuning on the UMPH cohort, additional modalities, such as laboratory tests, can be incorporated in the same manner. The resulting multimodal representations are concatenated and passed through a fusion multilayer perceptron (MLP) to construct longitudinal patient representations, which are subsequently combined with time-invariant features to predict pancreatic cancer risk (Fig. 3A). This framework is modular and can be readily extended to accommodate further modalities as they become available, such as in the refined model using UMPH data.

Diagnosis Prediction Intervals

Compared with prior studies[21,22], we adopted a more strict single-interval prediction design. When predicting "PANC diagnosis within N months", all EHR information on visits occurring within the N-month interval before the PANC diagnosis was masked and only earlier EHR information was retained to train the model (Fig. 3B). Specifically, for each PANC patient, we extracted all diagnoses, procedures, and medication records occurring prior to the PANC diagnosis date. We truncated the clinical trajectories according to different predefined prediction intervals N (N=6-, 12-, 24-, 36-, 60- month). For non-PANC controls, the reference date was defined as the end of their observation period, and their records were truncated backward from this time point using the same prediction intervals. This strategy is essential to reduce the risk of data leakage from prodromal symptoms very close to the diagnosis date and provides a more realistic assessment for early detection potential.

Implementation Details & Evaluation Metrics

Patients with valid longitudinal trajectories after preprocessing are randomly split into training, validation, and test sets in an 8:1:1 ratio. The model was trained using the binary cross-entropy loss with logits, optimized with AdamW (initial learning rate of 10^{-4} , weight decay applied). All training experiments were performed on 4 NVIDIA A40 GPUs for 10 epochs with a batch size of 128. Model performance was primarily assessed using several metrics: (1) the area under the receiver operating characteristic curve (AUROC), (2) the area under the precision-recall-gain curve (AUPRG),[26] normalized adjusted area under the precision–recall curve (AUPRC) with precision-gain and recall-gain scores, which handles the imbalances between case vs control samples. For each metric, we used the 95% confidence interval (CI) estimated via 200 bootstrap resamples. To provide a more comprehensive assessment, we also reported precision, recall, F1 score, and specificity (Supplementary Table 2-5).

Feature Interpretation

To enhance model interpretability, we applied integrated gradients (IG)[27] to perform attribution analysis on the input embeddings. IG is a path-integrated, gradient-based attribution method that quantifies feature contributions to the prediction outcome (see Supplementary Appendix for formulation). Specifically, it computes feature attributions by accumulating gradients of the prediction score with respect to the input along a straight-line path in the embedding space, interpolating from a baseline input to the actual input. We defined the baseline as a zero embedding vector, representing the absence of any clinical information. Based on the attribution results, we aggregated importance scores for all inputs and identified the clinical features most relevant to PANC risk prediction.

Results

Model Building on MarketScan Cohort

For each prediction interval setting, we obtained the model that achieved the minimum validation loss on the validation set, then we evaluated their predictive ability on the MarketScan hold-out test set. As shown in Figs. 4A-B, the 6-month prediction window achieves the best performance, with an AUROC of 0.812 (95% CI, 0.807–0.818) and an AUPRG of 0.851 (95% CI, 0.843–0.860). The AUROC and AUPRG decrease as the prediction period gets longer, as expected. Still, the 60-month prediction window reaches the AUROC of 0.735 (95% CI, 0.724–0.745) and AUPRG of 0.629 (95% CI, 0.600–0.656). The shorter prediction window also has larger training patient size, likely contributing to better model performance (Supplementary Table 6). To show where PANCDetect stands, we compared our model with CancerRiskNet[21], with the same data preprocessing and the same training, validation, and testing cohorts. PANCDetect consistently

outperforms CancerRiskNet across all prediction intervals (averaged AUC 0.633-0.765), particularly on long term predictions over 36 and 60 months (Supplementary Fig. 2).

We further evaluated several LLMs pretrained on biomedical corpora, such as BioGPT[17], which is trained on PubMed abstracts, and GatorTron[28], which is trained on large-scale de-identified clinical notes of University of Florida (UF) Health, PubMed articles, and Wikipedia. These domain-specific models were compared against the general-purpose embedding model OpenAI text-embedding-3-small, with predictive performance evaluated by AUROC and AUPRG on the test set. Experimental results show that the OpenAI text-embedding-3-small model achieves the performance that is comparable to, and in most cases better than, these specialized biomedical domain-specific models (Supplementary Table 7). This aligns with some of the existing literature showing that specialized models, whether trained from scratch or fine-tuned, often underperform compared to more general large-scale models across multiple tasks.[29,30]

We also performed another ablation study on multimodal fusion strategies by sequentially removing procedure and/or drug modalities (Supplementary Table 8). The model using diagnosis codes alone provided the baseline performance (AUROC of 0.779 and AUPRG of 0.810 for 6-month prediction). Adding procedures or drugs yielded consistent gains, with AUROC of 0.802 and AUPRG of 0.841, and AUROC of 0.801 and AUPRG of 0.828, for 6-month prediction respectively. The best results were obtained when all three modalities were combined (AUROC of 0.812 and AUPRG of 0.851 for 6-month prediction). The results show that incorporating both procedure and drug information led to a significant improvement on predictive accuracy, revealing the complementary nature of these modalities.

External Validation Results on UMPH and OneFlorida+ Cohorts

To further evaluate the generalizability of our models, we tested them on two additional EHR cohorts: the OneFlorida+ research network data and the UMPH data (Fig. 4C-F). The model

achieves very good predictions on OneFlorida+ cohort (Fig. 4C-D), particularly in short-term (6-12 months) prediction with AUROCs of 0.793 and 0.729, and AUPRGs of 0.917 and 0.859. The performance declines with longer exclusion intervals, as expected. On the other hand, for the UMPH Cohort, the model shows decent but lower predictive AUROC scores, ranging from 0.711 at 6-month prediction to 0.584 at 60-month prediction (Fig. 4E). AUPRG, however, decreases as the exclusion interval lengthens, from 0.891 at 6-month prediction to 0.598 at 60 months (Fig. 4F).

Fine-tuning results on UMPH dataset

The performance decline observed on the two EHR cohorts (particularly UMPH data), as compared to that of MarketScan hold-out data, is primarily attributable to the difference in EHR data structure as compared to the MarketScan claims-based data. This includes differences in coding systems, variable coverage, patterns of missingness,[31] as well as the underlying patient populations, age distributions, and care-seeking behaviors, etc. Collectively, these factors contribute to input feature distribution shifts that limit the generalizability of the models trained on the MarketScan cohort originally. Therefore, it is necessary to fine-tune the model on the EHR cohort to improve adaptability and healthcare-specific predictive performance.

Thus, we fine-tuned the model trained from the MarketScan cohort on the UMPH hospital EHR data, and incorporated laboratory test results as an additional Lab modality. We kept the data processing pipeline and model architecture unchanged. As shown in Fig. 4G-H, the fine-tuning results drastically improved the model's performance on UMPH test data. The AUROC reached 0.927 (95% CI, 0.899–0.951) and the AUPRG reached 0.979 (95% CI, 0.968–0.986) for the 6-month prediction window, whereas the 60-month prediction window achieved an AUROC of 0.835 (95% CI, 0.783–0.879) and an AUPRG of 0.911 (95% CI, 0.849–0.955).

Top risk predictor for pancreatic cancer

Based on IG attribution analysis, we summarized the top ten risk factors identified across different prediction windows (Fig.5). In the MarketScan cohort, the model consistently highlighted type 2 diabetes mellitus and other diseases of pancreas (except acute pancreatitis) as the top 1 and top 2 most important risk factors across all prediction windows, respectively. Even though their absolute scores decrease with the increase of prediction interval, these abnormalities maintained high importance scores throughout, indicating that they are the most stable and significant signals preceding pancreatic cancer.[32–34] Abdominal & pelvic pain and acute pancreatitis are additional top risk factors as well, though their importance decreases as the prediction interval increases. By 60-month prediction, abdominal & pelvic pain falls below top 10 risk factors.

In addition, personal or family history of malignant neoplasms repeatedly appear among the top 10 risk factors (Fig. 5), reflecting their importance in the preclinical phase of pancreatic cancer and potential comorbidity mechanism. In particular, neoplasm in different organs, such as colon and other GI tract organs, prostate and breast are also identified as important risk factors across different prediction windows, suggesting possible shared risk pathways or genetic susceptibility mechanisms with pancreatic cancer. Previous epidemiological studies have confirmed that these factors indeed elevate the risk of pancreatic cancer, thereby providing strong external evidence of our attribution result.[35–39]

Discussion

This study leveraged the large-scale MarketScan database to develop a predictive model using multi-modal EHR information, including diagnoses, procedures, and medications. Incorporating laboratory test data during fine-tuning on UM hospital validation cohort further improved model performance.

The strengths of this study lie in three areas. First, at the data level, the MarketScan database—covering a large and diverse population—provided a solid foundation for robust model development. Moreover, we show that inclusion of multimodal attributes in the EHR data can substantially improve the accuracy of cancer risk prediction through the ablation study. In particular, the inclusion of laboratory test results in hospital datasets significantly increased sensitivity and precision in the fine-tuned model. Second, at the methodological level, the proposed LLM-based embedding plus multimodal fusion framework also contributes to superior performance. This is demonstrated by the better performance of PANCDetect in comparison to CancerRiskNet, using the same training, validation, and testing cohort. Interestingly, our embedding comparison revealed that representations derived from general LLMs outperformed those trained exclusively on biomedical text, underscoring the cross-domain transferability and robustness of general-purpose models. Third, at the study design level, the integration of IG-based interpretability revealed clinically meaningful drivers of predictions.

Nevertheless, several limitations remain. First, the generalizability of the model requires further improvement, as performance may decline when applied to external institutions or different populations. Second, discrepancies in coding systems across healthcare settings have not been resolved, posing challenges for real-world deployment and cross-system adaptation. Thirdly, longer exclusion intervals tend to exclude a much larger number of patients and truncate the available trajectories, resulting in shorter and potentially less informative input sequences. This increases the difficulty of the prediction task over a long period of time. Federated learning frameworks could be utilized to facilitate more powerful long-time risk prediction, while preserving patient privacy, thereby improving generalizability[40]. Lastly, future efforts should be paid on testing PANCDetect in prospective studies, before being deployed in real-world clinical environments for clinical decision support.

In conclusion, we present a multimodal pancreatic cancer risk prediction framework based on large-scale EHR data, which demonstrated superior performance and interpretability compared with existing methods. Our model shows promise for potential applications in high-risk pancreatic cancer screening, early diagnosis, and personalized interventions. Ultimately, it may improve pancreatic cancer survival outcomes while reducing healthcare costs.

Data Availability

This study analyzed de-identified EHR and claims data from three sources: the MarketScan Research databases (licensed from Merative), the University of Michigan Precision Health (UMPH), and the OneFlorida+ Clinical Research Network. Due to licensing restrictions, institutional regulations, and patient privacy considerations, these data are not publicly available. Access to MarketScan requires a data use agreement and paid license from Merative. Access to UMPH and OneFlorida+ data may be provided to qualified researchers upon reasonable request to the corresponding author and is subject to approval by the respective institutional review boards and data custodians.

Acknowledgement

The study was not supported by any external funding. We would like to thank Patrick Brady and Institute for HealthCare Policy and Innovation (IHPI) of University of Michigan for helping us to access the MarketScan data. LXG would like to thank her father Jianwen Xia, who died from pancreatic cancer on 09/28/2024 for his eternal love.

Author Contribution

LXG envisioned and supervised this project. ZJ, XG and ZW built and tested the model, conducted the analysis and wrote the manuscript. QY and RY conducted validation analysis using

FloridaOne+ cohort. XY and XZ helped with analysis. XY helped to access UMPH data. All authors have read, revised and approved the final manuscript.

Competing Interest

The authors declare no conflict of interest.

Reference

1. Al-Shaheri FN, Alhamdani MSS, Bauer AS, et al. Blood biomarkers for differential diagnosis and early detection of pancreatic cancer. Cancer Treat Rev 2021;96:102193. DOI: 10.1016/j.ctrv.2021.102193

2. Pereira SP, Oldfield L, Ney A, et al. Early detection of pancreatic cancer. Lancet Gastroenterol Hepatol 2020;5(7):698–710. DOI: 10.1016/S2468-1253(19)30416-9

3. Huang J, Lok V, Ngai CH, et al. Worldwide Burden of, Risk Factors for, and Trends in Pancreatic Cancer. Gastroenterology 2021;160(3):744–54. DOI: 10.1053/j.gastro.2020.10.007

4. Yang J, Xu R, Wang C, Qiu J, Ren B, You L. Early screening and diagnosis strategies of pancreatic cancer: a comprehensive review. Cancer Commun 2021;41(12):1257–74. DOI: 10.1002/cac2.12204

5. Søreide K, Ismail W, Roalsø M, Ghotbi J, Zaharia C. Early Diagnosis of Pancreatic Cancer: Clinical Premonitions, Timely Precursor Detection and Increased Curative-Intent Surgery. Cancer Control 2023;30:10732748231154711. DOI: 10.1177/10732748231154711

6. Elbanna KY, Jang H-J, Kim TK. Imaging diagnosis and staging of pancreatic ductal adenocarcinoma: a comprehensive review. Insights Imaging 2020;11(1):58. DOI: 10.1186/s13244-020-00861-y

7. Wittram R, Kreis L, König H-H, Brettschneider C. Economic evaluations of early detection strategies for pancreatic cancer: a systematic review. Eur J Health Econ 2025. DOI: 10.1007/s10198-025-01793-4

8. Mishra AK, Chong B, Arunachalam SP, Oberg AL, Majumder S. Machine Learning Models for Pancreatic Cancer Risk Prediction Using Electronic Health Record Data—A Systematic Review and Assessment. Am J Gastroenterol 2024;119(8):1466–82. DOI: 10.14309/ajg.0000000000002870

9. Guo LL, Fries J, Steinberg E, et al. A multi-center study on the adaptability of a shared foundation model for electronic health records. Npj Digit Med 2024;7(1):171. DOI: 10.1038/s41746-024-01166-w

10. Xian S, Grabowska ME, Kullo IJ, et al. Transformer patient embedding using electronic health records enables patient stratification and progression analysis. Npj Digit Med 2025;8(1):521. DOI: 10.1038/s41746-025-01872-z

11. Zhang S, Cornet R, Benis N. Cross-Standard Health Data Harmonization using Semantics of Data Elements. Sci Data 2024;11(1):1407. DOI: 10.1038/s41597-024-04168-1

12. Li Y, Rao S, Solares JRA, et al. BEHRT: Transformer for Electronic Health Records. Sci Rep 2020;10(1):7155. DOI: 10.1038/s41598-020-62922-y

13. Li Y, Mamouei M, Salimi-Khorshidi G, et al. Hi-BEHRT: Hierarchical Transformer-Based Model for Accurate Prediction of Clinical Events Using Multimodal Longitudinal Electronic Health Records. IEEE J Biomed Health Inform 2023;27(2):1106–17. DOI: 10.1109/JBHI.2022.3224727

14. Rasmy L, Xiang Y, Xie Z, Tao C, Zhi D. Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. Npj Digit Med 2021;4(1):86. DOI: 10.1038/s41746-021-00455-y

15. Su X, Messica S, Huang Y, et al. Multimodal Medical Code Tokenizer. February 2, 2025 (https://arxiv.org/abs/2502.04397). Preprint.

16. Lee SA, Wu A, Chiang JN. Clinical ModernBERT: An efficient and long context encoder for biomedical text. April 4, 2025 (https://arxiv.org/abs/2504.03964). Preprint.

17. Luo R, Sun L, Xia Y, et al. BioGPT: Generative Pre-trained Transformer for Biomedical Text Generation and Mining. Brief Bioinform 2022;23(6):bbac409. DOI: 10.1093/bib/bbac409

18. Jiang LY, Liu XC, Nejatian NP, et al. Health system-scale language models are all-purpose prediction engines. Nature 2023;619(7969):357–62. DOI: 10.1038/s41586-023-06160-y

19. Jia K, Kundrot S, Palchuk MB, et al. A pancreatic cancer risk prediction model (Prism) developed and validated on large-scale US clinical data. eBioMedicine 2023;98:104888. DOI: 10.1016/j.ebiom.2023.104888

20. Appelbaum L, Cambronero JP, Stevens JP, et al. Development and validation of a pancreatic cancer risk model for the general population using electronic health records: An observational study. Eur J Cancer 2021;143:19–30. DOI: 10.1016/j.ejca.2020.10.019

21. Placido D, Yuan B, Hjaltelin JX, et al. A deep learning algorithm to predict risk of pancreatic cancer from disease trajectories. Nat Med 2023;29(5):1113–22. DOI: 10.1038/s41591-023-02332-5

22. Park J, Patterson J, Acitores Cortina JM, Gu T, Hur C, Tatonetti N. Enhancing EHR-based pancreatic cancer prediction with LLM-derived embeddings. Npj Digit Med 2025;8(1):465. DOI: 10.1038/s41746-025-01869-8

23. Butler AM, Nickel KB, Overman RA, Brookhart MA. IBM MarketScan Research Databases. In: Sturkenboom M, Schink T, eds. Databases for Pharmacoepidemiological Research. Cham: Springer International Publishing; 2021:243-251. DOI: 10.1007/978-3-030-51455-6_20

24. Zawistowski M, Fritsche LG, Pandit A, et al. The Michigan Genomics Initiative: A biobank linking genotypes and electronic clinical records in Michigan Medicine patients. Cell Genomics 2023;3(2):100257. DOI: 10.1016/j.xgen.2023.100257

25. Shenkman E, Hurt M, Hogan W, et al. OneFlorida Clinical Research Consortium: Linking a Clinical and Translational Science Institute With a Community-Based Distributive

Medical Education Model. Acad Med 2018;93(3):451–5. DOI: 10.1097/ACM.0000000000002029

26. Flach P, Kull M. Precision-Recall-Gain Curves: PR Analysis Done Right. Proc Neural Inf Process Syst. 2015:838-846. DOI: 10.5555/2969239.2969333

27. Sundararajan M, Taly A, Yan Q. Axiomatic Attribution for Deep Networks. Proc Int Conf Mach Learn. 2017:3319-3328. DOI: 10.5555/3305890.3306024

28. Yang X, Chen A, PourNejatian N, et al. A large language model for electronic health records. Npj Digit Med 2022;5(1):194. DOI: 10.1038/s41746-022-00742-2

29. Hegselmann S, von Arnim G, Rheude T, et al. Large Language Models are Powerful Electronic Health Record Encoders. February 24, 2025 (https://arxiv.org/abs/2502.17403). Preprint.

30. Excoffier J-B, Roehr T, Figueroa A, Papaioannou J-M, Bressem K, Ortala M. Generalist embedding models are better at short-context clinical semantic search than specialized embedding models January 3, 2024 (https://arxiv.org/abs/2401.01943). Preprint.

31. Laws MB, Michaud J, Shield R, McQuade W, Wilson IB. Comparison of Electronic Health Record–Based and Claims-Based Diabetes Care Quality Measures: Causes of Discrepancies. Health Serv Res 2018;53(S1):2988–3006. DOI: 10.1111/1475-6773.12819

32. Cho J, Scragg R, Petrov MS. Postpancreatitis Diabetes Confers Higher Risk for Pancreatic Cancer Than Type 2 Diabetes: Results From a Nationwide Cancer Registry. Diabetes Care 2020;43(9):2106–12. DOI: 10.2337/dc20-0207

33. Huxley R, Ansary-Moghaddam A, Berrington de González A, Barzi F, Woodward M. Type-II diabetes and pancreatic cancer: a meta-analysis of 36 studies. Br J Cancer 2005;92(11):2076–83. DOI: 10.1038/sj.bjc.6602619

34. Lai S-W, Liao K-F. Pancreatitis, Type 2 Diabetes Mellitus, and Pancreatic Cancer (Diabetes Metab J 2025;49:252-63). Diabetes Metab J 2025;49(3):518–9. DOI: 10.4093/dmj.2025.0237

35. Price SJ, Gibson N, Hamilton WT, King A, Shephard EA. Intra-abdominal cancer risk with abdominal pain: a prospective cohort primary care study. Br J Gen Pract J R Coll Gen Pract 2022;72(718):e361–8. DOI: 10.3399/BJGP.2021.0552

36. Matsubayashi H, Kiyozumi Y, Ishiwatari H, Uesaka K, Kikuyama M, Ono H. Surveillance of Individuals with a Family History of Pancreatic Cancer and Inherited Cancer Syndromes: A Strategy for Detecting Early Pancreatic Cancers. Diagn Basel Switz 2019;9(4):169. DOI: 10.3390/diagnostics9040169

37. Seppälä TT, Burkhart RA, Katona BW. Hereditary colorectal, gastric, and pancreatic cancer: comprehensive review. BJS Open 2023;7(3):zrad023. DOI: 10.1093/bjsopen/zrad023

38. Mocci E, Milne RL, Méndez-Villamil EY, et al. Risk of pancreatic cancer in breast cancer families from the breast cancer family registry. Cancer Epidemiol Biomark Prev Publ Am Assoc Cancer Res Cosponsored Am Soc Prev Oncol 2013;22(5):803–11. DOI: 10.1158/1055-9965.EPI-12-0195

39. Jacobs EJ, Chanock SJ, Fuchs CS, et al. Family history of cancer and risk of pancreatic cancer: a pooled analysis from the Pancreatic Cancer Cohort Consortium (PanScan). Int J Cancer 2010;127(6):1421–8. DOI: 10.1002/ijc.25148

40. Sadilek A, Liu L, Nguyen D, et al. Privacy-first health research with federated learning. Npj Digit Med 2021;4(1):132. DOI: 10.1038/s41746-021-00489-2

Figure Legend

Fig 1. Overall workflow of this study.

(1) Cohort selection: The MarketScan cohort was used for model training, while two independent cohorts (UMPH and OneFlorida+) served as external validation datasets. (2) Feature Engineering: During model training, diagnoses, procedures, and medications were extracted as inputs. Standardized textual descriptions of these medical codes were encoded into embeddings using a pretrained LLM. (3) Model Building and Validation: The generated embeddings were passed into a Transformer-based model, with the best model selected based on minimized validation loss. Evaluation was conducted on the held-out test set and external validation cohorts. Attribution analysis using Integrated Gradients (IG) was applied to identify key risk factors. (4) Model Refinement: Fine-tuning was further performed on the UMPH cohort to include additional laboratory tests into the pretrained model using MarketScan data. UMPH: University of Michigan Precision Health; PANC: pancreatic cancer.

Fig 2. Characteristics of the MarketScan, UMPH and OneFlorida+ cohorts.

Panels A,C, and E show the age-sex distribution of pancreatic cancer patients on MarketScan (Panel A), UMPH (Panel C) and OneFlorida+ (Panel E) cohorts, respectively. Panels B,D, and F show heatmaps of the joint distribution of diagnosis trajectory length and the number of distinct disease codes for PANC and non-PANC patients in the MarketScan (Panel B), UMPH (Panel D), and OneFlorida+ (Panel F) cohorts. Color intensity reflects the number of patients in each bin of age group, with darker shades indicating larger patient counts. UMPH: University of Michigan Precision Health; PANC: pancreatic cancer.

Fig 3, Model structure of PANCDetect and single-interval prediction design.

Panel A shows the overall model structure of PANCDetect. Firstly, time-varying EHR events (diagnoses, procedures, medications) are encoded with LLM-based embeddings and temporal markers to represent longitudinal patient trajectories and then processed by a Transformer with masked self- and cross-attention. The resulting sequence representation is pooled into a patient vector and late-fused by concatenation with time-invariant attributes (e.g., sex, geographic location); a final MLP maps the fused vector to a continuous pancreatic cancer risk score. Panel B illustrates the single-interval prediction design, where EHR events within *N* months before cancer diagnosis are masked to reduce data leakage, and trajectories are truncated according to predefined prediction windows (6, 12, 24, 36, and 60 months). LLM: large language model; MLP: multi-layer perceptron; PANC: pancreatic cancer.

Fig 4. Model performance across MarketScan, UMPH, and OneFlorida+ cohorts.

Panels A and B show results of the model trained on MarketScan training set tested on MarketScan hold-out test data; Panels C and D show OneFlorida+ external validation results using the model trained on MarketScan training data; Panels E and F show UMPH external validation results on the model trained on MarketScan training data; Panels G and H show fine-tuning results on UMPH hold-out data, after adding laboratory test data to the model built on MarketScan data. Panels A,C,E, and G: Results in AUROC; Panels B,D,F, and H: Results in AUPRG. Shaded areas represent 95% confidence intervals. AUROC: area under the receiver operating characteristic curve; AUPRG: area under the precision-recall-gain curve; UMPH: University of Michigan Precision Health; PANC: pancreatic cancer.

Fig. 5. Attribution analysis of top EHR risk predictors for pancreatic cancer.

Top 10 features from the model trained on MarketScan data. The top predictive clinical features across different prediction windows are ranked by the IG attribution scores.

Tables

Table 1: Summary of patient statistics in the three datasets of this study

General dataset information	**MarketScan**	**UMPH**	**OneFlorida+**
Dataset timeline	2009-2022	2006-2025	2013-2024
Total n patients	680,821	48,318	54,853
Male (%)	346,084(50.8%)	25,751(53.3%)	27,596 (50.3%)
Female (%)	334,737(49.2%)	22,567(46.7%)	27,257 (49.7%)
Median n disease codes per patient	101	298	77
Median length of trajectory in years	3	5	2
Pancreatic cancer cohort information			
Total n patients	113,518	8,257	10,500
Male (%)	57,713(50.8%)	4,406(53.4%)	5,537 (52.7%)
Female (%)	55,805(49.2%)	3,851(46.6%)	4,963 (47.3%)
Median age at PC diagnosis	62	67	67
Median n disease codes per patient	242	193	96
Median length of trajectory in years	3	1	1
Smoker (%)	24,965(22.0%)	891(10.8%)	4,319 (41.1%)
Alcohol user (%)	5,520（4.9%）	239(2.9%)	693 (6.6%)
Illegal drug user (%)	3,819（3.4%）	328(4.0%)	712 (6.8%)

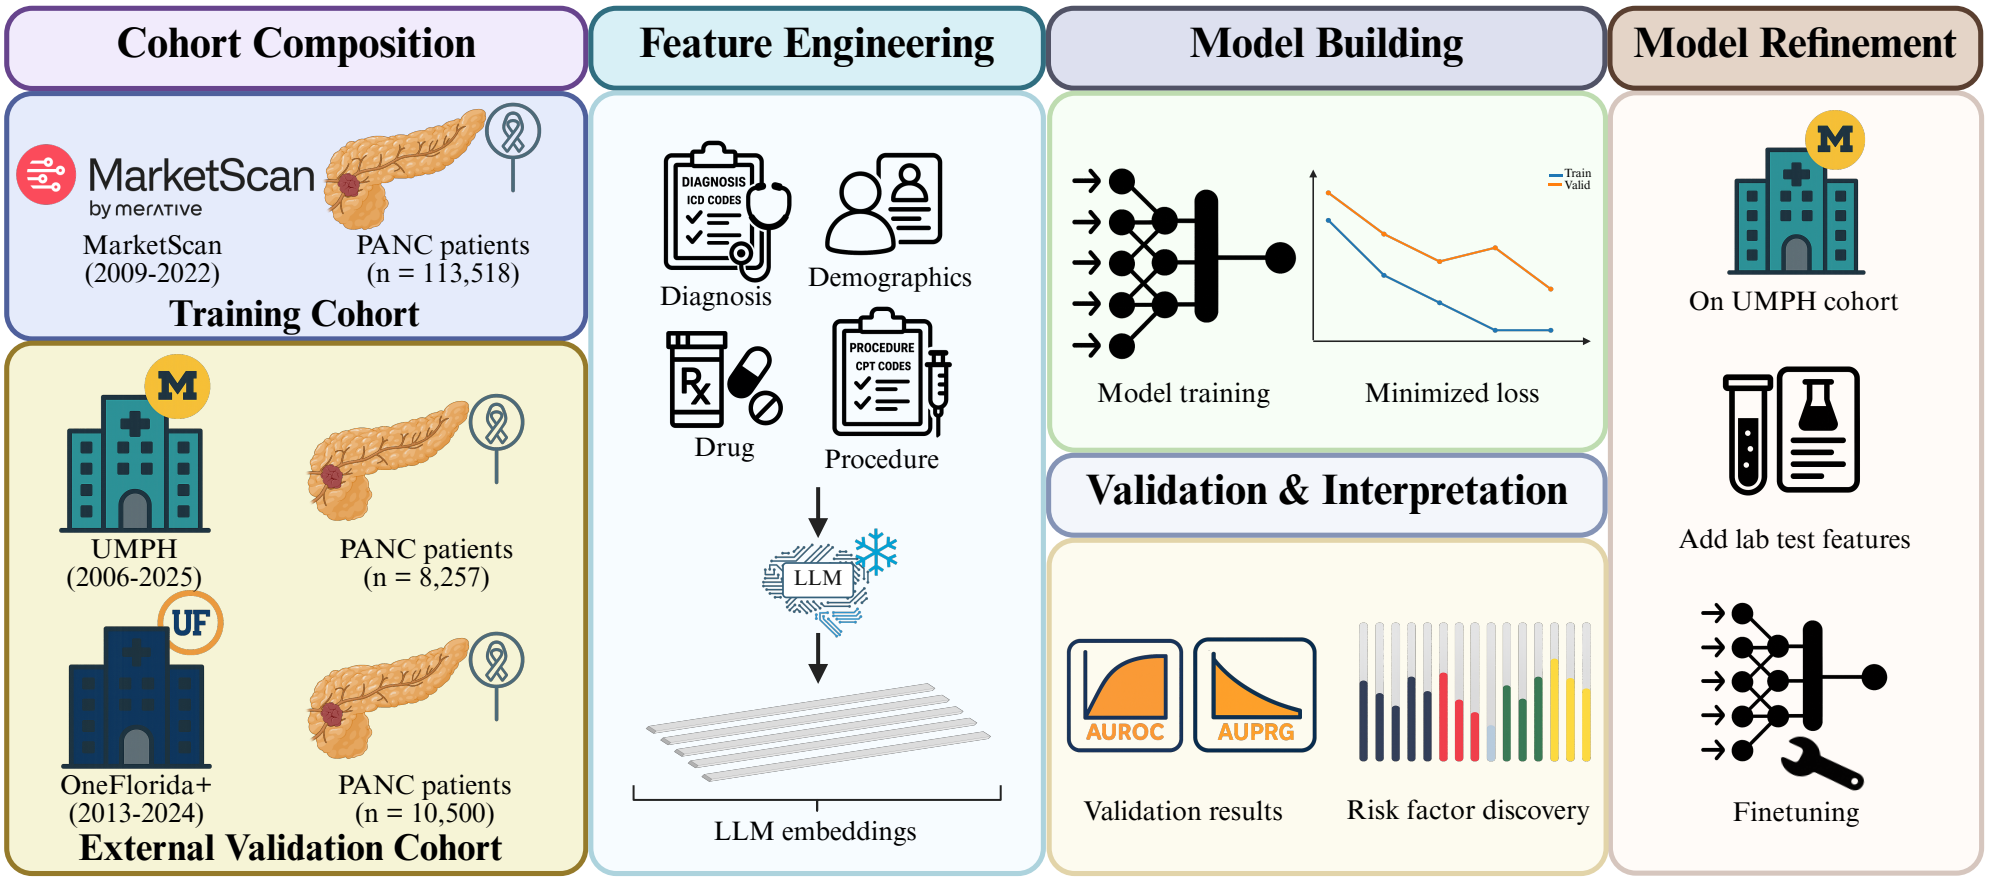
Cohort Composition
MarketScan
by merative
MarketScan
(2009-2022)
PANC patients
(n = 113,518)
Training Cohort
UMPH
(2006-2025)
PANC patients
(n = 8,257)
UF
OneFlorida+
(2013-2024)
PANC patients
(n = 10,500)
External Validation Cohort
Feature Engineering
DIAGNOSIS
ICD CODES
Diagnosis
Demographics
Drug
PROCEDURE
CPT CODES
Procedure
LLM
LLM embeddings
Model Building
Model training
Train
Valid
Minimized loss
Validation & Interpretation
AUROC
AUPRG
Validation results
Risk factor discovery
Model Refinement
On UMPH cohort
Add lab test features
Finetuning

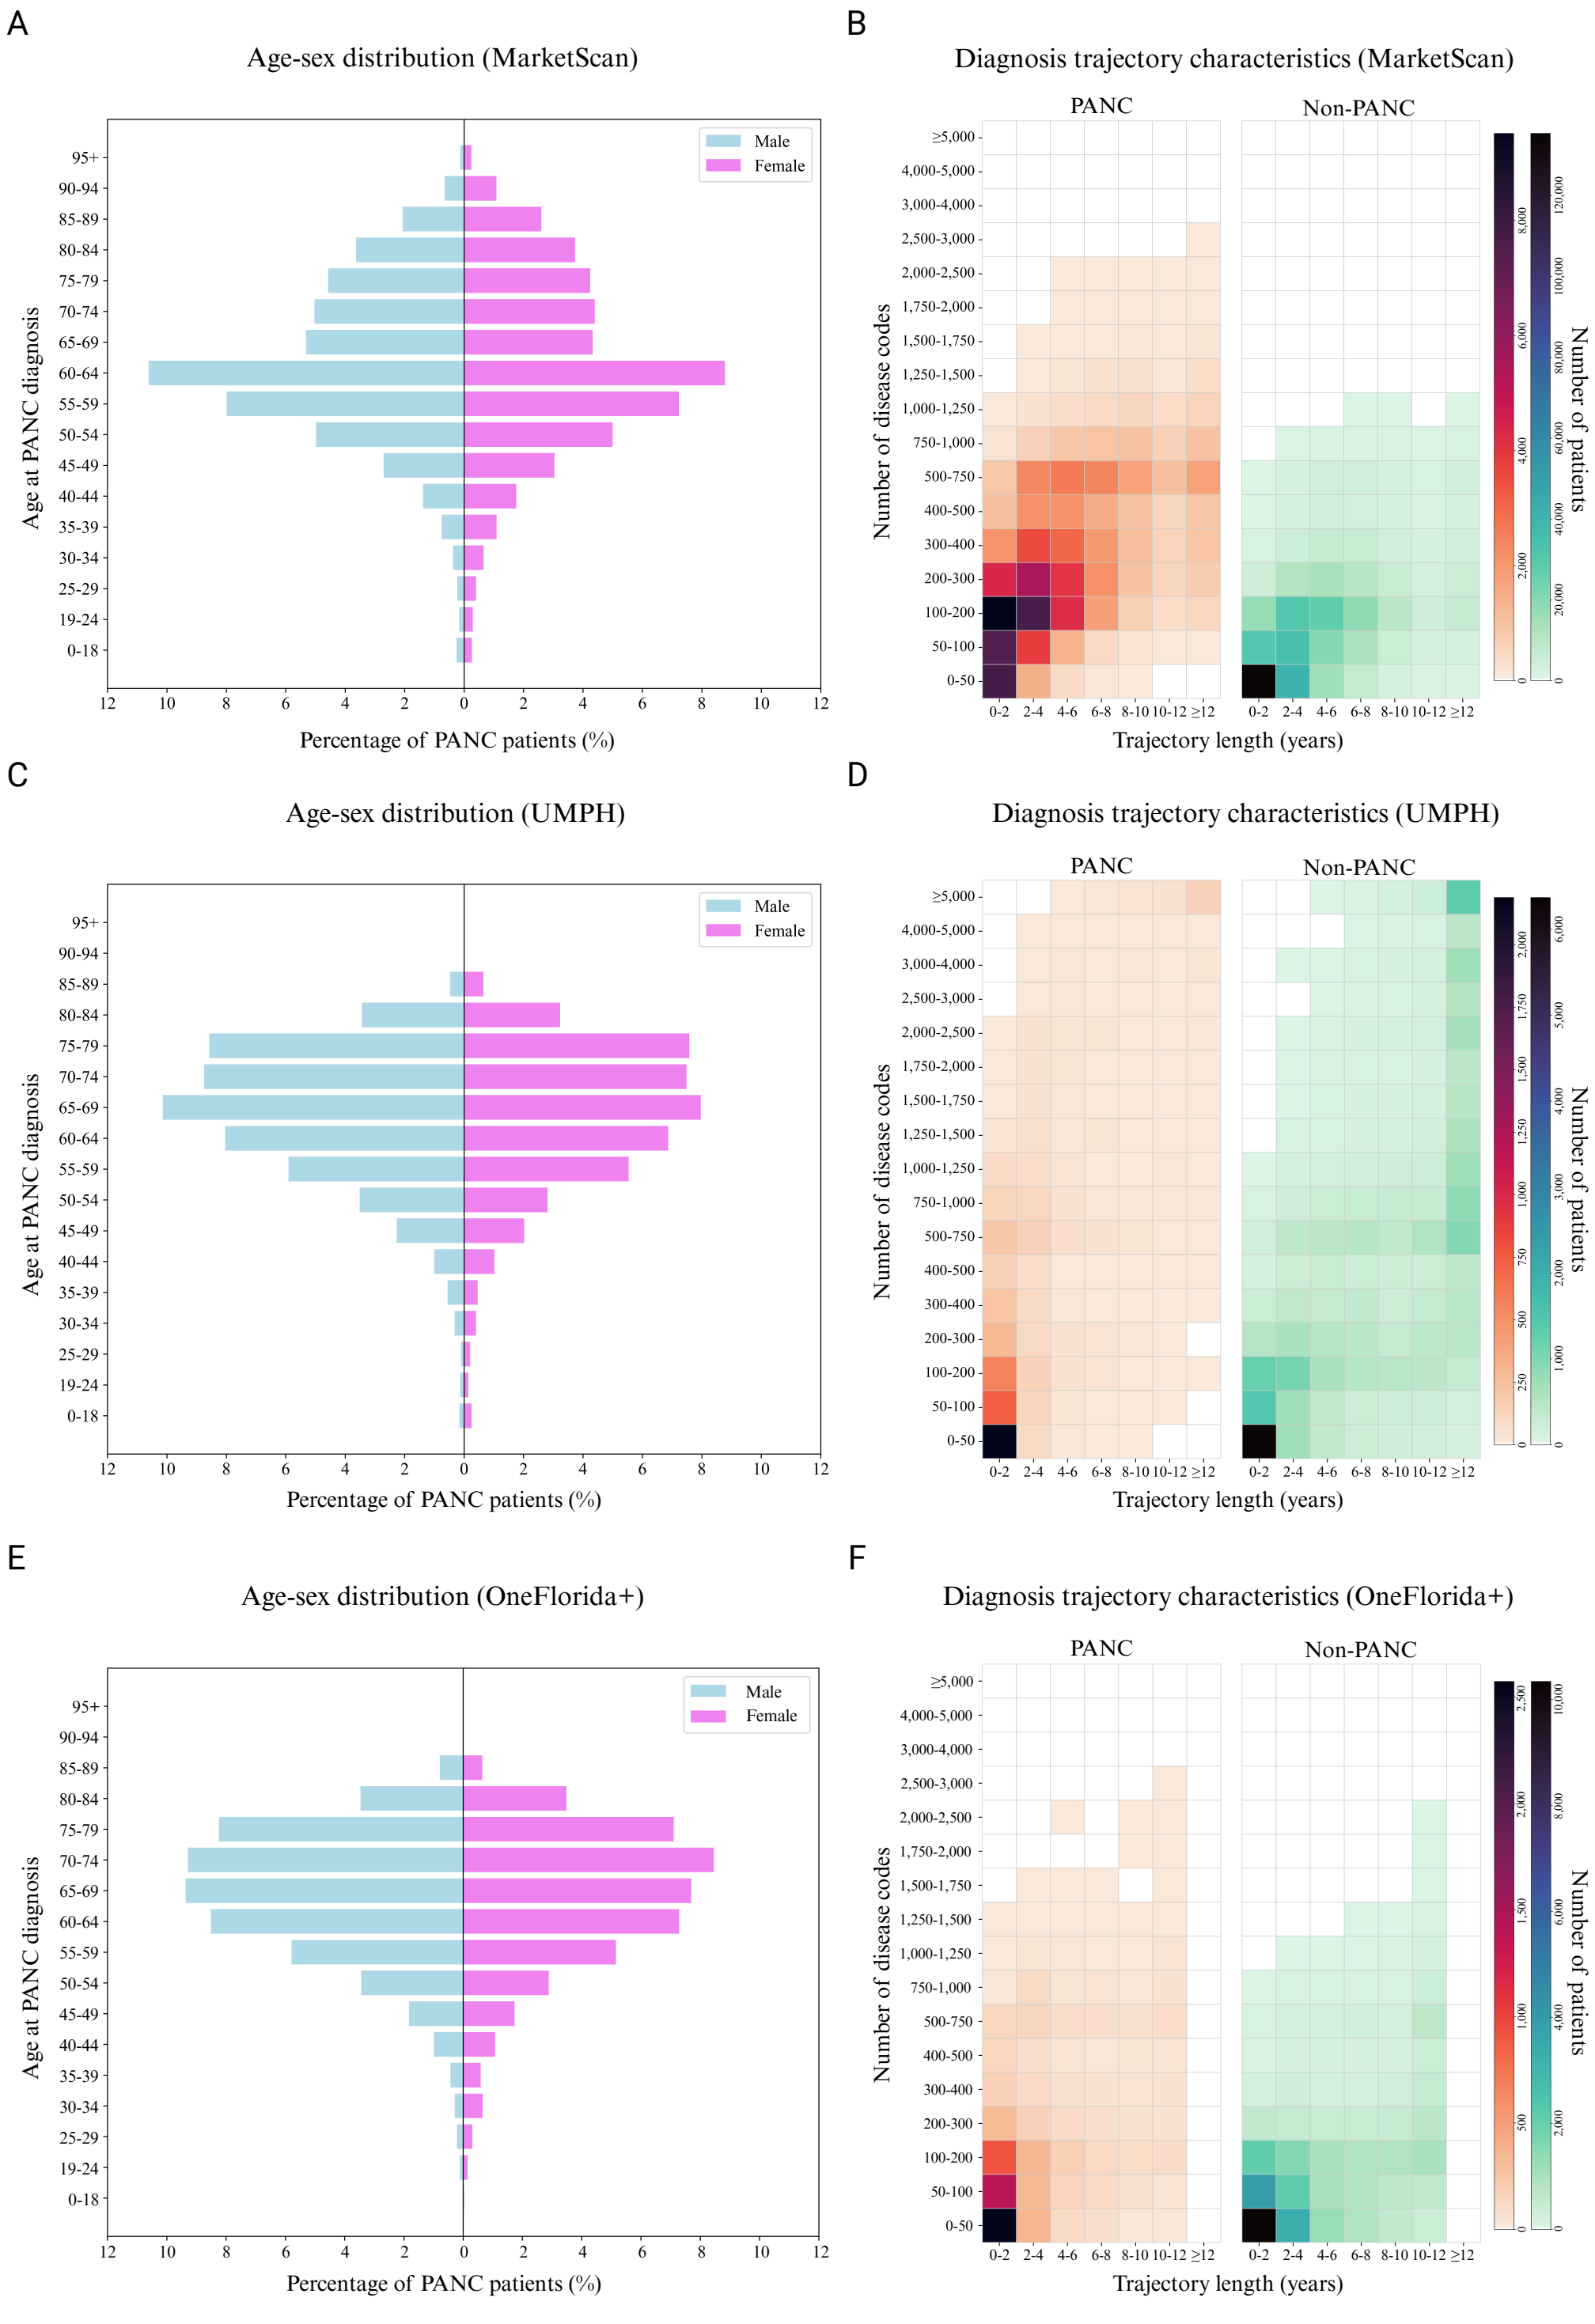
A
Age-sex distribution (MarketScan)
Male
Female
Age at PANC diagnosis
95+
90-94
85-89
80-84
75-79
70-74
65-69
60-64
55-59
50-54
45-49
40-44
35-39
30-34
25-29
19-24
0-18
12 10 8 6 4 2 0 2 4 6 8 10 12
Percentage of PANC patients (%)
B
Diagnosis trajectory characteristics (MarketScan)
PANC
Non-PANC
Number of disease codes
≥5,000
4,000-5,000
3,000-4,000
2,500-3,000
2,000-2,500
1,750-2,000
1,500-1,750
1,250-1,500
1,000-1,250
750-1,000
500-750
400-500
300-400
200-300
100-200
50-100
0-50
0-2 2-4 4-6 6-8 8-10 10-12 ≥12
Trajectory length (years)
Number of patients
C
Age-sex distribution (UMPH)
Male
Female
Age at PANC diagnosis
Percentage of PANC patients (%)
D
Diagnosis trajectory characteristics (UMPH)
PANC
Non-PANC
Number of disease codes
Trajectory length (years)
Number of patients
E
Age-sex distribution (OneFlorida+)
Male
Female
Age at PANC diagnosis
Percentage of PANC patients (%)
F
Diagnosis trajectory characteristics (OneFlorida+)
PANC
Non-PANC
Number of disease codes
Trajectory length (years)
Number of patients

A Overall model architecture of PANCDetect

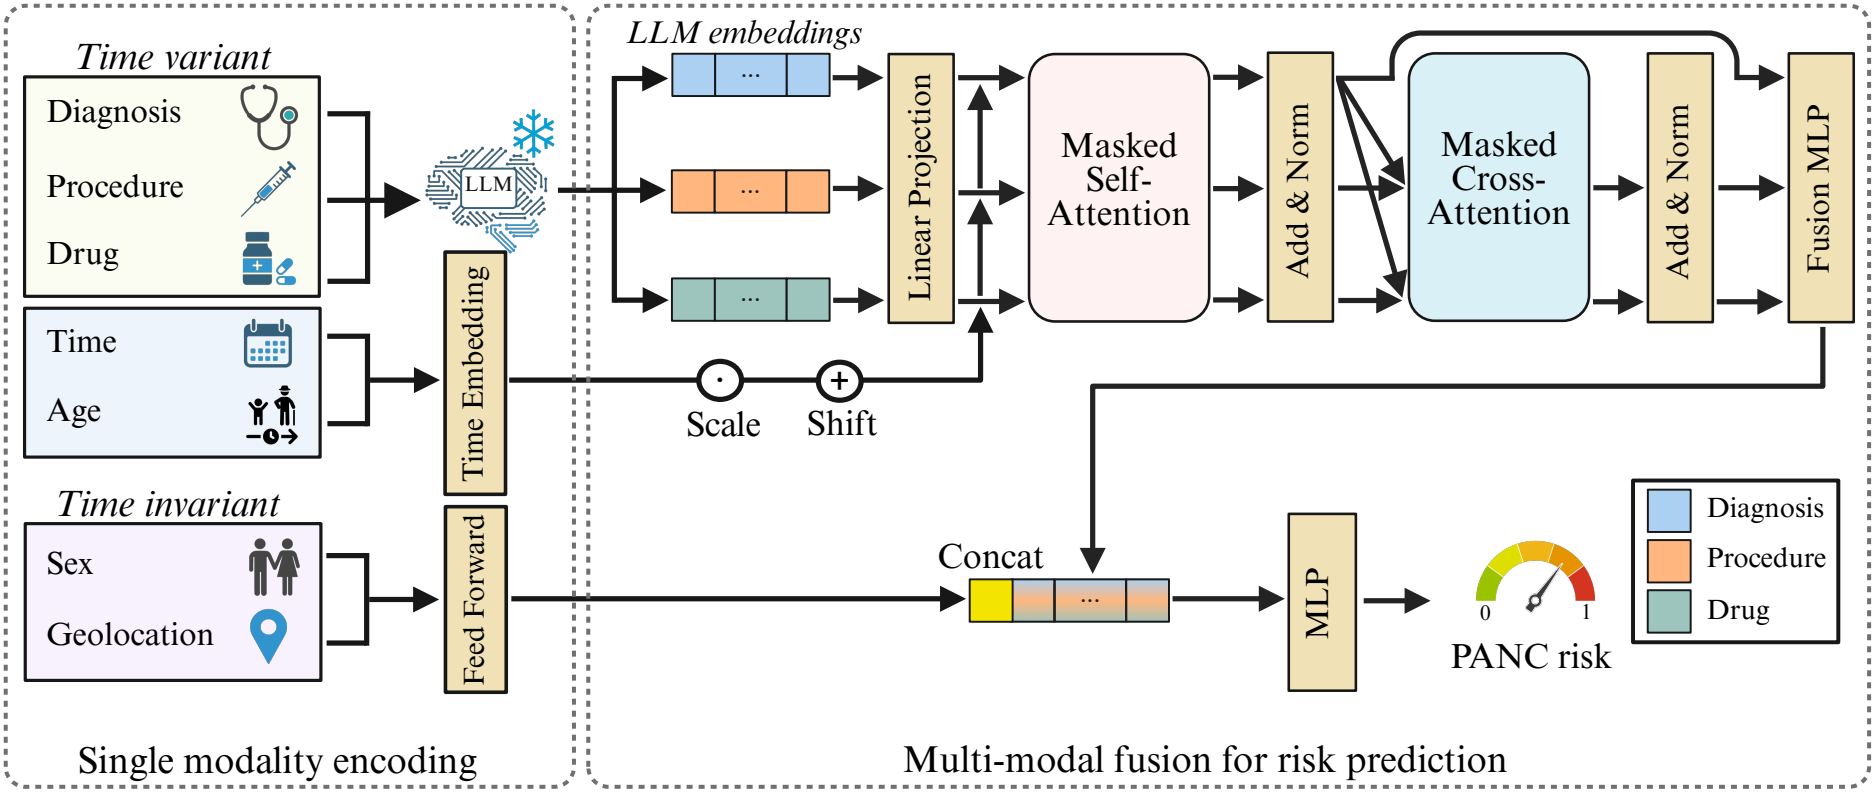

B Timepoints and prediction intervals

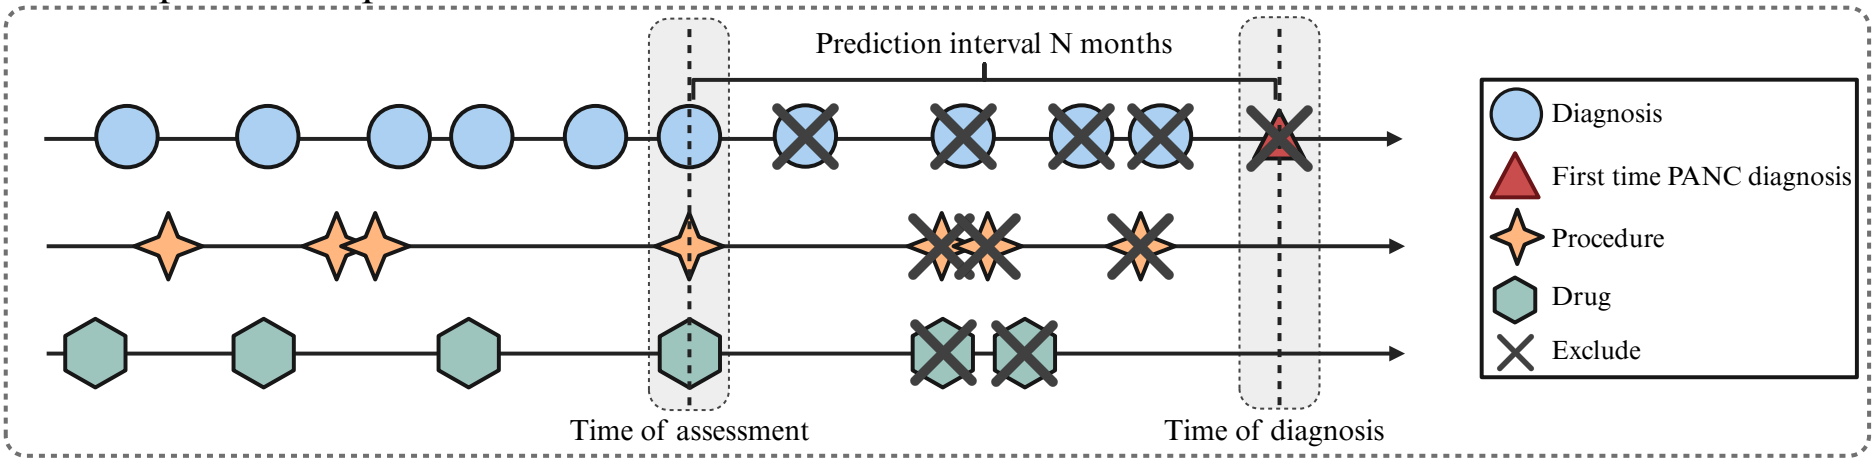

Test results on MarketScan holdout test set

A AUROC

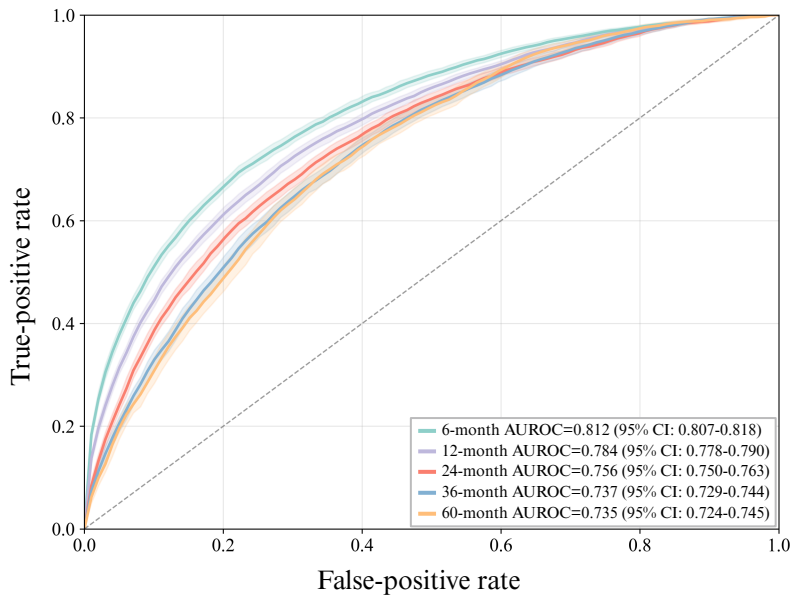

B AUPRG

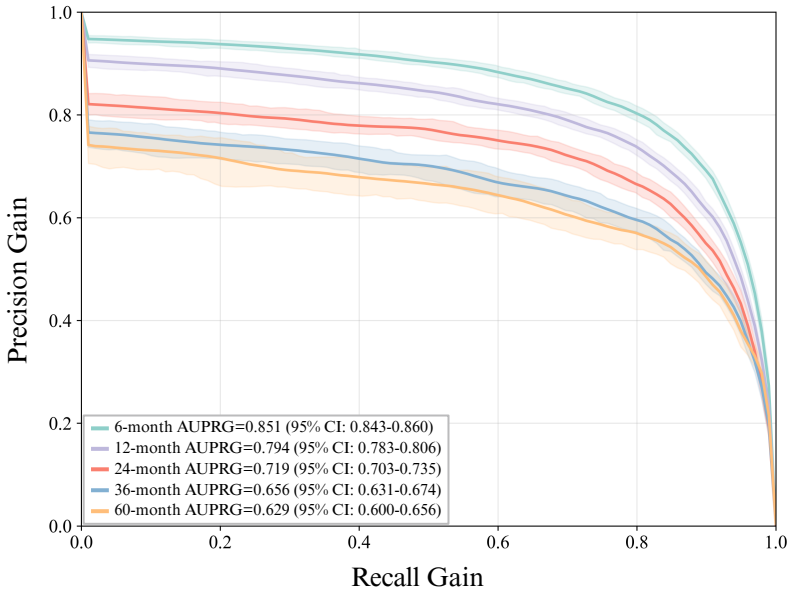

Validation results on OneFlorida+ cohort

C AUROC

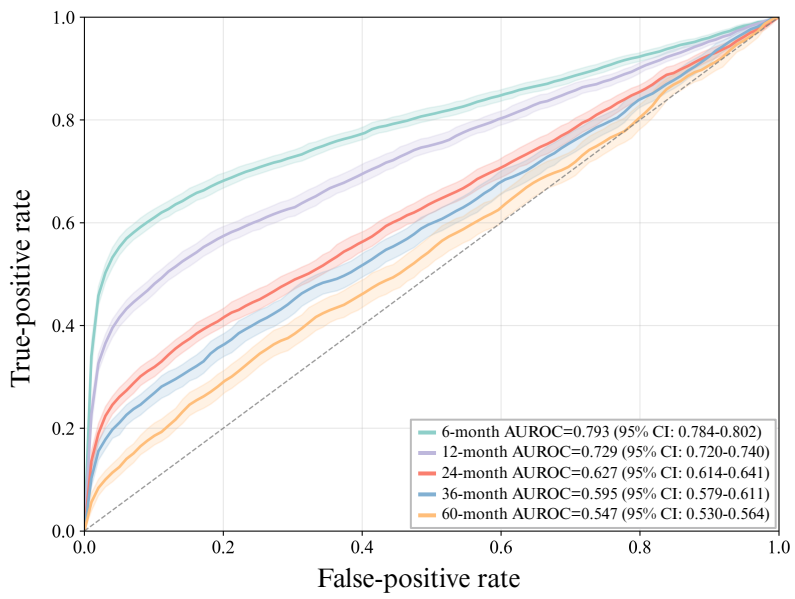

D AUPRG

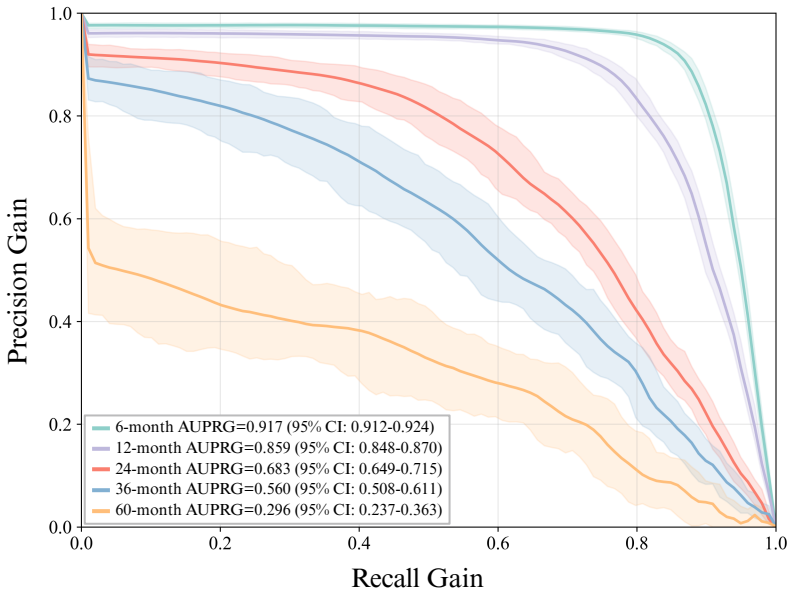

Validation results on UMPH cohort

E AUROC

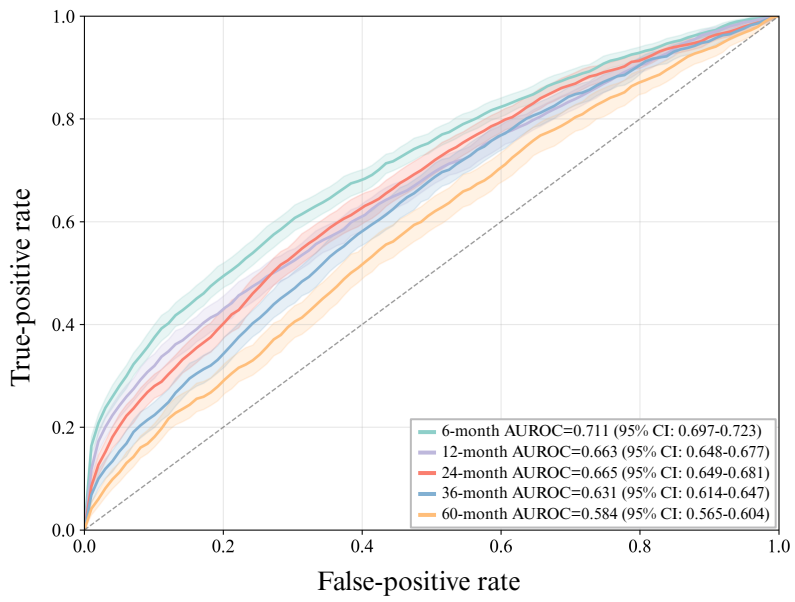

F AUPRG

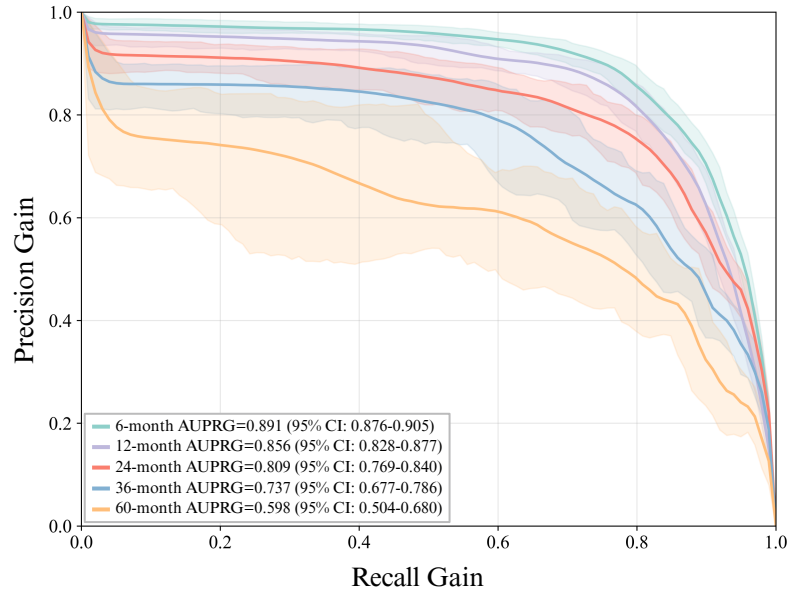

Finetuing results on UMPH holdout test set

G AUROC

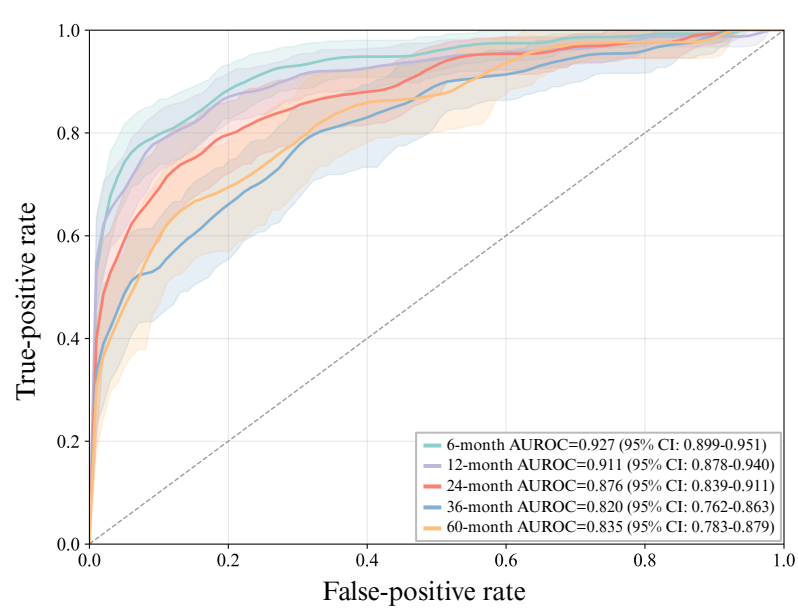

H AUPRG

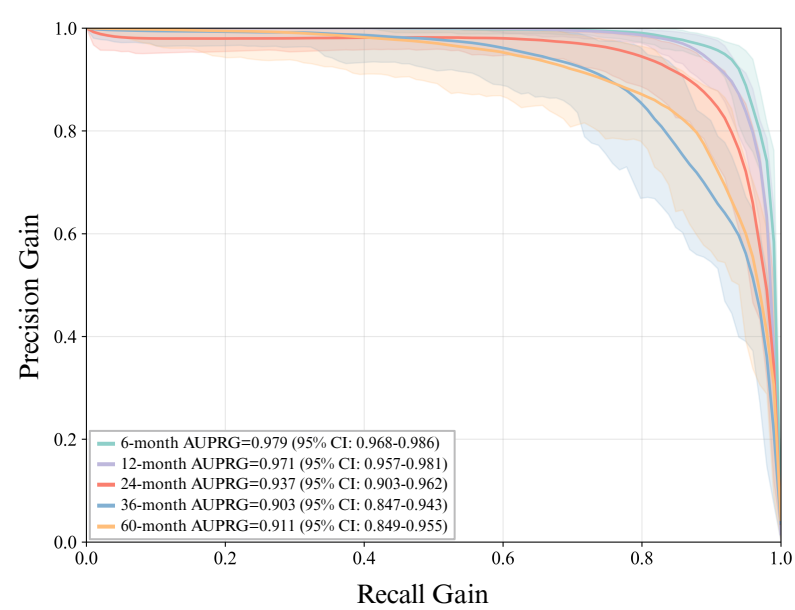

Feature attribution results

6-month Prediction	Score	12-month Prediction	Score	24-month Prediction	Score	36-month Prediction	Score	60-month Prediction	Score
Type 2 diabetes mellitus	2.070	Type 2 diabetes mellitus	1.118	Type 2 diabetes mellitus	0.777	Type 2 diabetes mellitus	0.358	Type 2 diabetes mellitus	0.259
Other diseases of pancreas	1.169	Other diseases of pancreas	0.749	Other diseases of pancreas	0.395	Other diseases of pancreas	0.183	Other diseases of pancreas	0.081
Abdominal and pelvic pain	0.773	Abdominal and pelvic pain	0.214	Benign neoplasm of colon, rectum, anus and anal canal	0.230	Benign neoplasm of colon, rectum, anus and anal canal	0.087	Benign neoplasm of colon, rectum, anus and anal canal	0.041
Acute pancreatitis	0.444	Acute pancreatitis	0.207	Abdominal and pelvic pain	0.207	Family history of malignant neoplasm	0.082	Family history of malignant neoplasm	0.034
Malignant neoplasm of prostate	0.277	Malignant neoplasm of prostate	0.157	Personal history of malignant neoplasm	0.182	Abdominal and pelvic pain	0.070	Personal history of malignant neoplasm	0.027
Benign neoplasm of colon, rectum, anus and anal canal	0.258	Personal history of malignant neoplasm	0.138	Acute pancreatitis	0.163	Acute pancreatitis	0.063	Malignant neoplasm of prostate	0.026
Malignant neoplasm of colon	0.216	Benign neoplasm of other digestive organs	0.106	Malignant neoplasm of prostate	0.158	Personal history of malignant neoplasm	0.061	Malignant neoplasm of kidney, except renal pelvis	0.024
Personal history of malignant neoplasm	0.198	Malignant neoplasm of colon	0.101	Malignant neoplasm of breast	0.098	Malignant neoplasm of breast	0.052	Acute pancreatitis	0.023
Benign neoplasm of other digestive organs	0.178	Genetic susceptibility to malignant neoplasm	0.100	Malignant neoplasm of colon	0.090	Malignant neoplasm of prostate	0.044	Other disorders of kidney and ureter	0.022
Other diseases of biliary tract	0.153	Family history of malignant neoplasm	0.098	Family history of malignant neoplasm	0.087	Benign neoplasm of other digestive organs	0.040	Malignant neoplasm of breast	0.019

ICD-10 chapters

- II Neoplasms
- IV Endocrine, nutritional and metabolic diseases
- XI Diseases of the digestive system
- XIV Diseases of the genitourinary system
- XVIII Symptoms, signs and abnormal clinical and laboratory findings, not elsewhere classified
- XXI Factors influencing health status and contact with health services

Supplementary Appendix

Integrated gradient (IG)[1] implementation details:

The IG along the ith dimension is defined as the following:

$$IG_i(x) := (x_i - x_i') \times \int_0^1 \frac{\partial F(x' + \alpha(x - x'))}{\partial x_i} d\alpha$$

where x is the actual input, x' is the baseline (often chosen as zeros or some "neutral" input), F is the prediction function (e.g., the model's risk score for pancreatic cancer), α is an interpolation parameter that smoothly moves from baseline (α=0) to actual input (α=1).

For implementation, the integral can't be computed analytically for deep models, we approximate the integration via the summation with steps m=30 as the following:

$$IG_i^{approx}(x) := (x_i - x_i') \times \sum_{k=1}^{m} \frac{\partial F(x' + \frac{k}{m}(x - x'))}{\partial x_i} \times \frac{1}{m}$$

So practically, IG is "average gradient × input difference."

Supplementary Fig. 1: Age and sex distributions of PANC patients (left) and matched non-PANC controls (right) in the MarketScan cohort. A 1:5 case–control matching strategy based on sex, age, and geographic location ensured consistent demographic ratios between cases and controls. We kept the same matching strategy for UMPH and OneFlorida+ cohorts too.

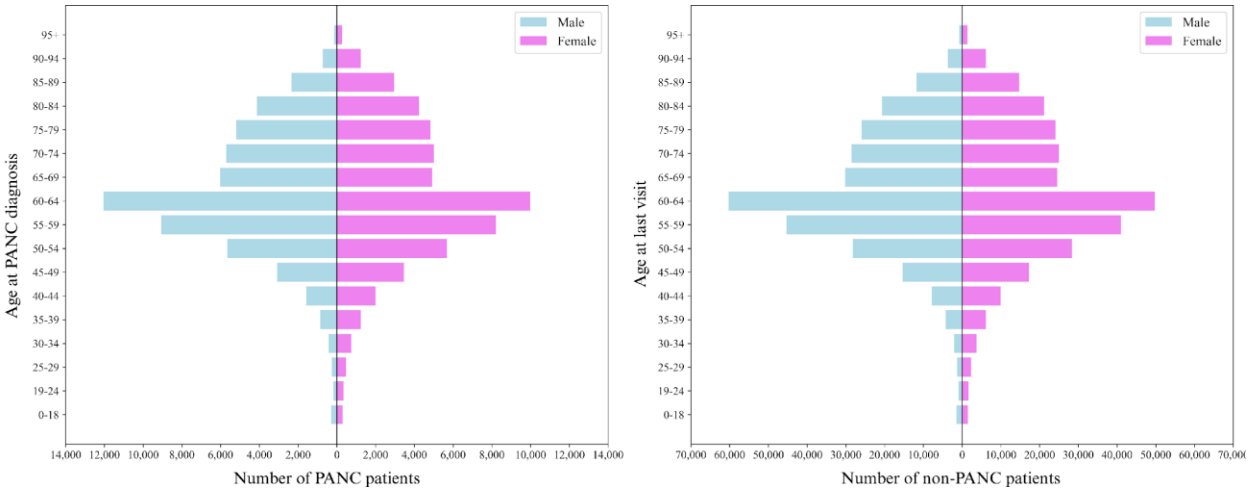

Supplementary Fig. 2: Comparative performance of PANCDetect and CancerRiskNet under the same prediction interval setting. PANCDetect achieved higher AUROC (A) and AUPRG (B) scores across all intervals, with particularly strong gains in long-term predictions (36- and 60-month). CancerRiskNet showed lower discrimination and precision–recall gain, especially at extended horizons. Shaded areas represent 95% confidence intervals.

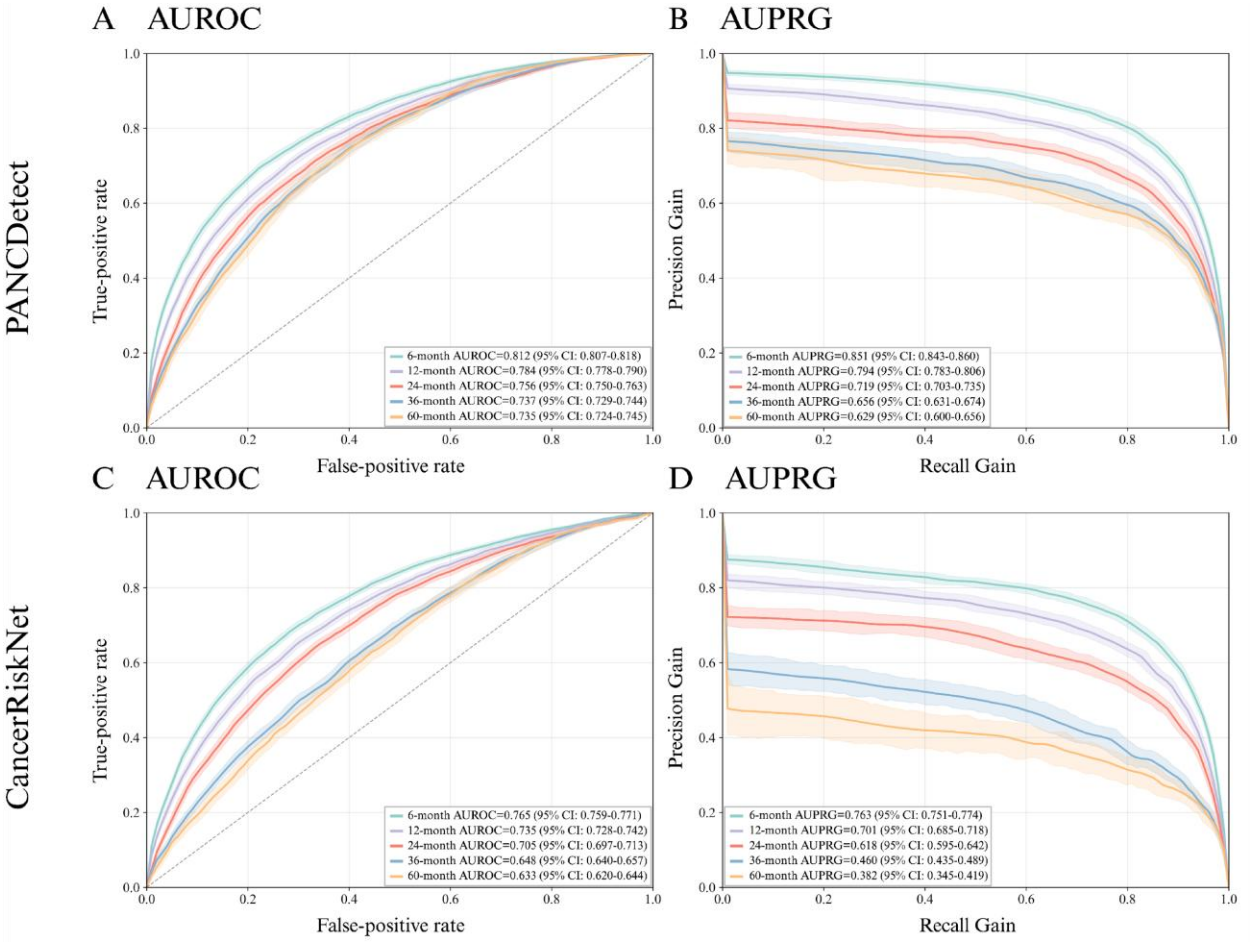

Supplementary Table 1: ICD 9/10 Code for pancreatic cancer

ICD 9 Code	ICD 10 Code
157: Malignant neoplasm of pancreas	C25: Malignant neoplasm of pancreas
157.0: Malignant neoplasm of head of pancreas	C25.0: Malignant neoplasm of head of pancreas
157.1: Malignant neoplasm of body of pancreas	C25.1: Malignant neoplasm of body of pancreas
157.2: Malignant neoplasm of tail of pancreas	C25.2: Malignant neoplasm of tail of pancreas
157.3: Malignant neoplasm of pancreatic duct	C25.3: Malignant neoplasm of pancreatic duct
157.4: Malignant neoplasm of islets of langerhans	C25.4: Malignant neoplasm of endocrine pancreas
157.8: Malignant neoplasm of other specified sites of pancreas	C25.7: Malignant neoplasm of other parts of pancreas
157.9: Malignant neoplasm of pancreas, part unspecified	C25.8: Malignant neoplasm of overlapping sites of pancreas
	C25.9: Malignant neoplasm of pancreas, unspecified

Supplementary Table 2: Other evaluation metrics performance on MarketScan test set.

	6-month	**12-month**	**24-month**	**36-month**	**60-month**
Precision	0.626 (95% CI: 0.612-0.638)	0.634 (95% CI: 0.616-0.653)	0.502 (95% CI: 0.484-0.525)	0.575 (95% CI: 0.537-0.607)	0.433 (95% CI: 0.405-0.455)
Recall	0.360 (95% CI: 0.350-0.369)	0.235 (95% CI: 0.225-0.246)	0.258 (95% CI: 0.246-0.270)	0.106 (95% CI: 0.097-0.116)	0.306 (95% CI: 0.290-0.324)
F1 Score	0.458 (95% CI: 0.445-0.467)	0.343 (95% CI: 0.332-0.357)	0.341 (95% CI: 0.327-0.355)	0.180 (95% CI: 0.165-0.194)	0.359 (95% CI: 0.340-0.377)
Specificity	0.955 (95% CI: 0.953-0.957)	0.970 (95% CI: 0.969-0.972)	0.944 (95% CI: 0.941-0.947)	0.982 (95% CI: 0.980-0.984)	0.901 (95% CI: 0.895-0.907)

Supplementary Table 3: Other evaluation metrics performance on UMPH cohort (external validation).

	6-month	**12-month**	**24-month**	**36-month**	**60-month**
Precision	0.200 (95% CI: 0.189-0.214)	0.212 (95% CI: 0.196-0.233)	0.154 (95% CI: 0.137-0.169)	0.132 (95% CI: 0.116-0.152)	0.063 (95% CI: 0.055-0.071)
Recall	0.373 (95% CI: 0.350-0.391)	0.244 (95% CI: 0.224-0.262)	0.225 (95% CI: 0.201-0.245)	0.130 (95% CI: 0.110-0.149)	0.223 (95% CI: 0.197-0.254)
F1 Score	0.260 (95% CI: 0.244-0.275)	0.227 (95% CI: 0.211-0.245)	0.183 (95% CI: 0.166-0.199)	0.131 (95% CI: 0.113-0.152)	0.098 (95% CI: 0.086-0.111)
Specificity	0.900 (95% CI: 0.897-0.904)	0.948 (95% CI: 0.945-0.950)	0.938 (95% CI: 0.935-0.941)	0.961 (95% CI: 0.959-0.964)	0.875 (95% CI: 0.871-0.879)

Supplementary Table 4: Other evaluation metrics performance on OneFlorida+ cohort (external validation).

	6-month	**12-month**	**24-month**	**36-month**	**60-month**
Precision	0.482 (95% CI: 0.467-0.493)	0.502 (95% CI: 0.484-0.521)	0.351 (95% CI: 0.331-0.371)	0.447 (95% CI: 0.409-0.480)	0.272 (95% CI: 0.239-0.303)
Recall	0.618 (95% CI: 0.600-0.634)	0.448 (95% CI: 0.429-0.464)	0.319 (95% CI: 0.299-0.335)	0.203 (95% CI: 0.185-0.224)	0.170 (95% CI: 0.147-0.196)
F1 Score	0.541 (95% CI: 0.527-0.553)	0.473 (95% CI: 0.457-0.488)	0.334 (95% CI: 0.317-0.351)	0.279 (95% CI: 0.253-0.302)	0.209 (95% CI: 0.182-0.234)
Specificity	0.891 (95% CI: 0.887-0.896)	0.927 (95% CI: 0.924-0.931)	0.898 (95% CI: 0.894-0.903)	0.955 (95% CI: 0.950-0.958)	0.915 (95% CI: 0.908-0.922)

Supplementary Table 5: Other evaluation metrics performance on UMPH hold-out test cohort (finetuning).

	6-month	**12-month**	**24-month**	**36-month**	**60-month**
Precision	0.799 (95% CI: 0.736-0.862)	0.809 (95% CI: 0.750-0.872)	0.736 (95% CI: 0.649-0.819)	0.807 (95% CI: 0.695-0.906)	0.673 (95% CI: 0.541-0.780)
Recall	0.686 (95% CI: 0.610-0.756)	0.654 (95% CI: 0.583-0.767)	0.565 (95% CI: 0.496-0.641)	0.400 (95% CI: 0.333-0.481)	0.438 (95% CI: 0.337-0.547)
F1 Score	0.738 (95% CI: 0.679-0.793)	0.724 (95% CI: 0.664-0.775)	0.639 (95% CI: 0.575-0.700)	0.535 (95% CI: 0.457-0.617)	0.530 (95% CI: 0.417-0.632)
Specificity	0.969 (95% CI: 0.958-0.979)	0.967 (95% CI: 0.956-0.980)	0.957 (95% CI: 0.941-0.973)	0.980 (95% CI: 0.968-0.991)	0.957 (95% CI: 0.939-0.973)

Supplementary Table 6: Cohort size settings for model training, validation, and finetuning.

	Case/Control				
Cohorts	**6-month**	**12-month**	**24-month**	**36-month**	**60-month**
MarketScan (Model building)					
Training	59,106/283,37	50,499/238,61	38,387/175,82	29,022/129,31	16,557/70,00
Validation	7,485/35,324	6,393/29,747	4,796/21,980	3,652/16,140	2,111/8,709
Hold-out testing	7,429/35,381	6,401/29,739	4,832/21,945	3,720/16,073	2144/8,676
OneFlorida+ (External validation)					
Testing	3,878/23,743	3,074/18,784	2,297/13,307	1,815/10,051	1,187/6,309
UMPH (External validation)					
Testing	1,758/26,326	1,585/27,416	1,334/26,728	1,129/24,829	804/21,252
UMPH (Finetuning)					
Training	1,422/6,822	1,262/6,166	1,072/5,192	902/4,416	633/3,146
Validation	180/851	161/768	124/659	112/553	91/381
Testing	156/875	162/767	138/646	115/550	80/393

Supplementary Table 7: Ablation study of different embedding generation methods.

Model name & embedding dimension	**6-month prediction**		**60-month prediction**	
	AUROC	**AUPRG**	**AUROC**	**AUPRG**
BioGPT (1024)	0.801 (95% CI: 0.793-0.805)	0.837 (95% CI: 0.829-0.846)	0.710 (95% CI: 0.698-0.721)	0.578 (95% CI: 0.547-0.607)
GatorTron (1024)	0.812 (95% CI: 0.806-0.817)	**0.853** (95% CI: 0.836-0.861)	0.729 (95% CI: 0.718-0.739)	0.623 (95% CI: 0.595-0.651)
Trained from scratch (128)	0.772 (95% CI: 0.764-0.777)	0.802 (95% CI: 0.791-0.812)	0.682 (95% CI: 0.670-0.694)	0.520 (95% CI: 0.484-0.553)
OpenAI text-embedding-3-small (1536)	**0.812** (95% CI: 0.807-0.818)	0.851 (95% CI: 0.843-0.860)	**0.735** (95% CI: 0.724-0.745)	**0.629** (95% CI: 0.600-0.656)

Supplementary Table 8: Ablation study of different modalities used for training.

Modalities	**6-month prediction**		**60-month prediction**	
	AUROC	**AUPRG**	**AUROC**	**AUPRG**
Diagnosis	0.779 (95% CI: 0.773-0.784)	0.810 (95% CI: 0.799-0.819)	0.708 (95% CI: 0.698-0.719)	0.558 (95% CI: 0.527-0.586)
Diagnosis + Procedure	0.802 (95% CI: 0.796-0.807)	0.841 (95% CI: 0.832-0.850)	0.721 (95% CI: 0.711-0.732)	0.603 (95% CI: 0.575-0.627)
Diagnosis + Drug	0.801 (95% CI: 0.794-0.806)	0.828 (95% CI: 0.818-0.838)	0.729 (95% CI: 0.718-0.739)	0.618 (95% CI: 0.590-0.643)
Diagnosis + Procedure + Drug	**0.812** (95% CI: 0.807-0.818)	**0.851** (95% CI: 0.843-0.860)	**0.735** (95% CI: 0.724-0.745)	**0.629** (95% CI: 0.600-0.656)

Reference

1. Sundararajan M, Taly A, Yan Q. Axiomatic Attribution for Deep Networks. Proc Int Conf Mach Learn. 2017:3319-3328. DOI: 10.5555/3305890.3306024