

PanFoMa: A Lightweight Foundation Model and Benchmark for Pan-Cancer

Xiaoshui Huang^{1*}, Tianlin Zhu^{2*}, Yifan Zuo^{2†}, Xue Xia², Zonghan Wu³
Jiebin Yan², Dingli Hua², Zongyi Xu⁴, Yuming Fang², Jian Zhang⁵

¹Shanghai Jiao Tong University

²Jiangxi University of Finance and Economics

³East China Normal University

⁴Chongqing University of Posts and Telecommunications

⁵University of Technology Sydney

Abstract

Single-cell RNA sequencing (scRNA-seq) is essential for decoding tumor heterogeneity. However, pan-cancer research still faces two key challenges: learning discriminative and efficient single-cell representations, and establishing a comprehensive evaluation benchmark. In this paper, we introduce PanFoMa, a lightweight hybrid neural network that combines the strengths of Transformers and state-space models to achieve a balance between performance and efficiency. PanFoMa consists of a front-end local-context encoder with shared self-attention layers to capture complex, order-independent gene interactions; and a back-end global sequential feature decoder that efficiently integrates global context using a linear-time state-space model. This modular design preserves the expressive power of Transformers while leveraging the scalability of Mamba to enable transcriptome modeling, effectively capturing both local and global regulatory signals. To enable robust evaluation, we also construct a large-scale pan-cancer single-cell benchmark, PanFoMaBench, containing over 3.5 million high-quality cells across 34 cancer subtypes, curated through a rigorous pre-processing pipeline. Experimental results show that PanFoMa outperforms state-of-the-art models on our pan-cancer benchmark (+4.0%) and across multiple public tasks, including cell type annotation (+7.4%), batch integration (+4.0%) and multi-omics integration (+3.1%).

Introduction

The revolutionary advances in single-cell RNA sequencing (scRNA-seq) technology have provided an unprecedentedly powerful tool for systematically dissecting the heterogeneity of complex biological systems, such as tumors, at single-cell resolution (Jovic et al. 2022). By precisely capturing the gene expression profile of each cell, we can gain deep insights into the underlying mechanisms of tumor initiation, progression, metastasis, and response to therapy. Therefore, developing computational models capable of learning effective representations of cells and genes from high-dimensional, sparse transcriptomic data has become a

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

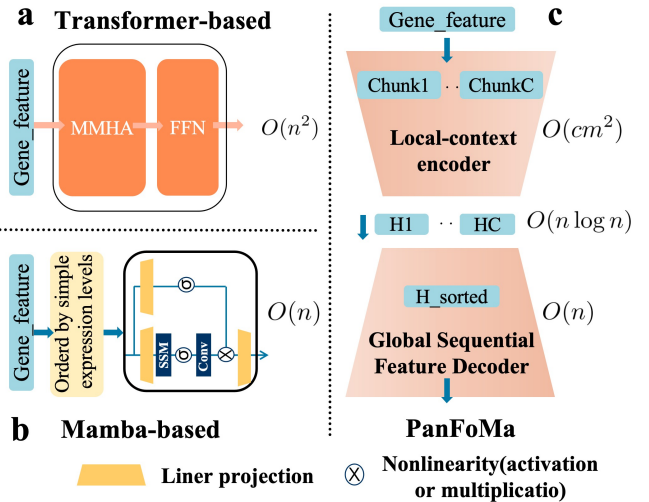


Figure 1: Comparison of architectures for modeling single-cell gene expression. (a) Transformer-based methods capture local features but incur $O(n^2)$ complexity. (b) Mamba-based models offer $O(n)$ efficiency but require fixed input ordering. (c) Our proposed PanFoMa captures both local and global dependencies via a hybrid encoder-decoder design, resulting in an overall computational complexity of $O(C \cdot M^2 + N \log N)$, where $N = C \cdot M$.

central challenge in computational biology. The deep representations learned by such models are foundational not only for advancing precision medicine and personalized diagnostics (Dutta et al. 2022), but also for a broad range of applications, including biomarker discovery, drug target identification, and fundamental studies of cellular processes (Van de Sande et al. 2023).

Most existing single-cell foundation models are transformer-based (He et al. 2024; Cui et al. 2024a; Hao et al. 2024; Theus et al. 2024; Cui et al. 2024b; Adduri et al. 2025; Fang et al. 2025; Zeng et al. 2025), with scGPT (Cui et al. 2024a) being a representative example. Inspired by advances in natural language processing, these models conceptualize genes as "tokens" and cells as "sentences," leveraging the self-attention mechanism to capture complex

dependencies among genes. Through generative pretraining on massive unlabeled single-cell datasets, they aim to learn rich and nuanced representations of the "gene language." However, these models face several inherent limitations. First, the *computational complexity* of the self-attention mechanism scales quadratically with the number of input genes, making it computationally prohibitive to process the complex transcriptome, which can contain tens of thousands of genes. Consequently, current approaches (Cui et al. 2024a,b) typically process only a subset of genes (e.g., 2048), selected via top-K highly variable genes (HVGs) (Cui et al. 2024a; Zeng et al. 2025), to reduce computational costs—at the expense of capturing only localized gene interactions. Second, this HVGs-based gene selection strategy (Cui et al. 2024a; Zeng et al. 2025) has notable drawbacks: it may *exclude important low-expression functional genes* (e.g., transcription factors) and introduces *analytical bias* into downstream tasks. These limitations compromise the models' ability to generalize across diverse biological contexts and hinder their interpretability in functional genomics applications.

To overcome the limitations of above transformer-based methods, recently, researchers have begun exploring the state-space models (Qi et al. 2025) for learning single-cell representations. Mamba (Gu and Dao 2023) offers linear-time complexity by maintaining a latent state that sequentially aggregates input information, thus potentially capturing global dependencies more efficiently than Transformers. However, directly applying Mamba to single-cell data presents several critical challenges. First, Mamba is fundamentally designed for ordered sequences, whereas a cell's gene expression profile is naturally an unordered set—gene interactions do not depend on any inherent sequence. To apply Mamba, existing approaches must impose a heuristic and fixed ordering of genes (e.g., by mean expression), thereby forcing an artificial sequence structure. *This static ordering is suboptimal*, as it neglects the context-dependent nature of gene functions, which emerge from complex and dynamic regulatory networks rather than simple expression levels. As a result, such ordering may impair the model's ability to capture authentic biological relationships. Second, although Mamba's latent state enables information integration across input sequences, its fixed dimensionality suffers from forgetting distant input tokens, which can *limit the ability to retain long-range dependencies*, particularly when dealing with long gene sequences. This constraint poses a serious challenge in tasks like Pan-Cancer classification and analysis, where capturing broad, global patterns across the transcriptome is essential. The limited memory capacity may result in degraded performance when modeling distant gene interactions critical for accurate diagnosis and interpretation.

To overcome the above limitations and achieve both high accuracy and computational efficiency in Pan-Cancer analysis, we propose a novel computational paradigm based on a decoupled modeling strategy for single-cell transcriptomes. Specifically, we separate the tasks of capturing local gene interactions and integrating global information, and assign them to dedicated architectures. A lightweight Transformer

encoder captures local, context-rich representations from a limited gene subset, while a Mamba-based decoder integrates these into a global cell-state representation using its efficient latent-state mechanism. This division of labor combines the expressive power of Transformers with the linear scalability of Mamba, resolving the trade-off between performance and efficiency. Biologically, it mimics the information flow from complex local regulatory signals to a unified cell state, enabling end-to-end modeling of the transcriptome.

We implement this idea in a innovative hybrid neural network named PanFoMa, which consists of two core modules. First, the **Local-context Encoder** partitions the input of thousands of genes into multiple fixed-size chunks and processes these chunks in parallel using a lightweight Transformer with shared parameters across layers. This design efficiently captures the complex interaction relationships within each chunk while significantly reducing computational and memory overhead. Second, the **Global Sequential Feature Decoder** introduces a novel processing pipeline: it first synthesizes a global vector representing the overall cell state by average pooling the summary information (`cls_token`) from all chunks. Next, it leverages this global vector to dynamically compute an "importance" score for each gene via dot product and intelligently sorts all genes accordingly. Finally, this biologically meaningful, ordered sequence is fed into a bidirectional Mamba module for deep integration, and the output is produced through a gated mechanism.

In addition, existing pan-cancer datasets (Weinstein et al. 2013; Štancl and Karlić 2023; Nofech-Mozes et al. 2023) often suffer from limited subtype coverage and sample size, introducing bias and overfitting in model evaluation. To address this, we construct and release a large-scale, high-quality benchmark dataset comprising over 3.5 million single cells across 34 major cancer types, following a rigorous data curation pipeline. This challenging benchmark requires strong generalization and is designed to drive progress in pan-cancer foundation model development.

In summary, the contributions are three aspects:

- We design a hierarchical *hybrid architecture* that integrates the strengths of Transformer-based local modeling and Mamba-based global integration, achieving a superior balance between performance and computational scalability.
- We develop a *global-informed dynamic sorting mechanism* that adaptively determines gene input sequences based on each cell's global transcriptomic context, moving beyond fixed heuristic rules and enhancing biological relevance.
- We construct a comprehensive *pan-cancer benchmark*, and through extensive evaluations on both curated and public benchmarks, we demonstrate that our model consistently outperforms existing state-of-the-art approaches, providing a robust framework to accelerate advances in precision medicine.

Related Work

Single cell foundation models

In recent years, drawing inspiration from successful paradigms in Natural Language Processing (NLP), Foundation Models have been widely applied to single-cell transcriptomics, aiming to learn *universal* cell and gene representations through large-scale self-supervised pre-training. The development in this field has primarily evolved along two core architectural lines: early models based on Transformers (Cui et al. 2024b,a; Zeng et al. 2025; Fang et al. 2025; Huang et al. 2024, 2025), and a new generation of architectures based on State Space Models (SSMs) (Qi et al. 2025) that have emerged in pursuit of higher efficiency.

The main idea of transformer-based methods leverages its powerful self-attention mechanism to capture complex, long-range interactions between genes. scGPT (Cui et al. 2024a) is one of the representative works in this area. It adopts a GPT-like generative pre-training framework (Radford et al. 2018; Zheng et al. 2024; Mei et al. 2022), designed to complete the entire gene expression profile based on partial gene or cell prompts. To address the fundamental challenge of the unordered nature of gene expression data, scGPT designed a special attention mask mechanism, which simulates an iterative generation process by randomly partitioning known and unknown genes during training, thereby enabling autoregressive-style learning on non-sequential data. In contrast, GeneFormer (Cui et al. 2024b) introduces a rank-based input encoding and masked gene prediction strategy to mitigate batch effects and better capture gene regulatory networks. He et al. (He et al. 2024) further fine-tuned GeneFormer for cancer applications by extending its pretraining with approximately 14 million cancer-derived cells. scFoundation (Hao et al. 2024) adopts an asymmetric transformer-like architecture to model complex gene relationships across diverse cell types. CancerFoundation (Theus et al. 2024) focuses exclusively on malignant cells during training. Most recently, CellFM (Fang et al. 2025) presents a single-cell foundation model with 800 million parameters, trained on 100 million human cells using a modified RetNet framework to balance efficiency and performance.

Despite the success of above Transformer-based models, they share an unavoidable bottleneck: the quadratic computational complexity, excluding low-expression function genes and analytical bias. To break through the first bottleneck, Mamba (Gu and Dao 2023) has provided a new solution for efficiently processing sequence data with linear computational complexity by using a latent state to accumulate the previous input tokens and propagate to the next token prediction. GeneMamba (Qi et al. 2025) is a Mamba-based foundation model specifically designed for single-cell transcriptomics. It inherits the expression-based ranking input strategy from GeneFormer but replaces the core computational backbone from a Transformer to a bidirectional Mamba (Bi-Mamba). Although Mamba-based architectures achieve a breakthrough in computational efficiency, their performance is highly dependent on a fixed, heuristic rule for creating a serialized input. However, this static sorting

mechanism is suboptimal because it overlooks a core biological fact: a gene's importance is not static but is determined by its dynamic functional role within a specific cellular context. Therefore, the field currently lacks a model that can simultaneously achieve a deep understanding of unordered gene relationships and highly efficient serialized processing.

Pan-cancer prediction

Pan-cancer prediction is a key area in cancer research, aiming to leverage machine learning and deep learning to identify shared and distinct features across different cancer types for improved diagnosis, prognosis, and treatment (Nopour 2024; Yang et al. 2025). Zhang et al. (Zhang et al. 2025) proposed a unified multimodal framework for pan-cancer survival prediction by integrating whole-slide histopathology images and gene expression data, addressing the generalization gap across cancer types. Bjerregaard-Michelsen et al. (Bjerregaard-Michelsen et al. 2025) compared pan-cancer and single-cancer models for short-term mortality prediction, highlighting the utility of ML in treatment decisions but with limited long-term forecasting power. Ocasio et al. (Ocasio et al. 2024) developed a transcriptome-based deep learning model for pan-cancer classification using TCGA data, though limited by dataset bias. Lareau et al. (Lareau, Parker, and Satpathy 2021) showed that integrating multiple datasets improves drug sensitivity prediction, while challenges remain in modeling drug resistance. Salmanpour et al. (Salmanpour et al. 2025) evaluates 226 studies applying machine learning to PET and SPECT imaging for cancer outcome prediction, revealing that deep radiomics and fusion models yield the best performance, with PET-based approaches outperforming SPECT. Ouyang et al. (Ouyang et al. 2025) reviewed existing models for drug response prediction, pointing out current limitations and the need for more practical guidance on model development. However, the current existing methods face challenges in multimodal data integration and model generalization (Singh et al. 2025). At the same time, it is crucial to address data bias in the training dataset and benchmarks by using diverse datasets for training and validation, thereby enhancing the generalizability of the models. This paper focuses on these challenges by proposing a foundation model and a novel pan-cancer benchmark.

Method

Overview

To address the dual challenges of accurate and efficient modeling whole-transcriptome single-cell data, we propose PanFoMa, a hybrid neural network based on a hierarchical local-to-global processing paradigm. The core idea is to decompose the transcriptome into manageable local blocks for parallel contextual encoding, followed by global integration via dynamic gene reordering. By unifying localized deep modeling and globally informed integration, PanFoMa provides a scalable and effective solution for high-resolution single-cell representation learning. The overall framework is illustrated in Figure 2.

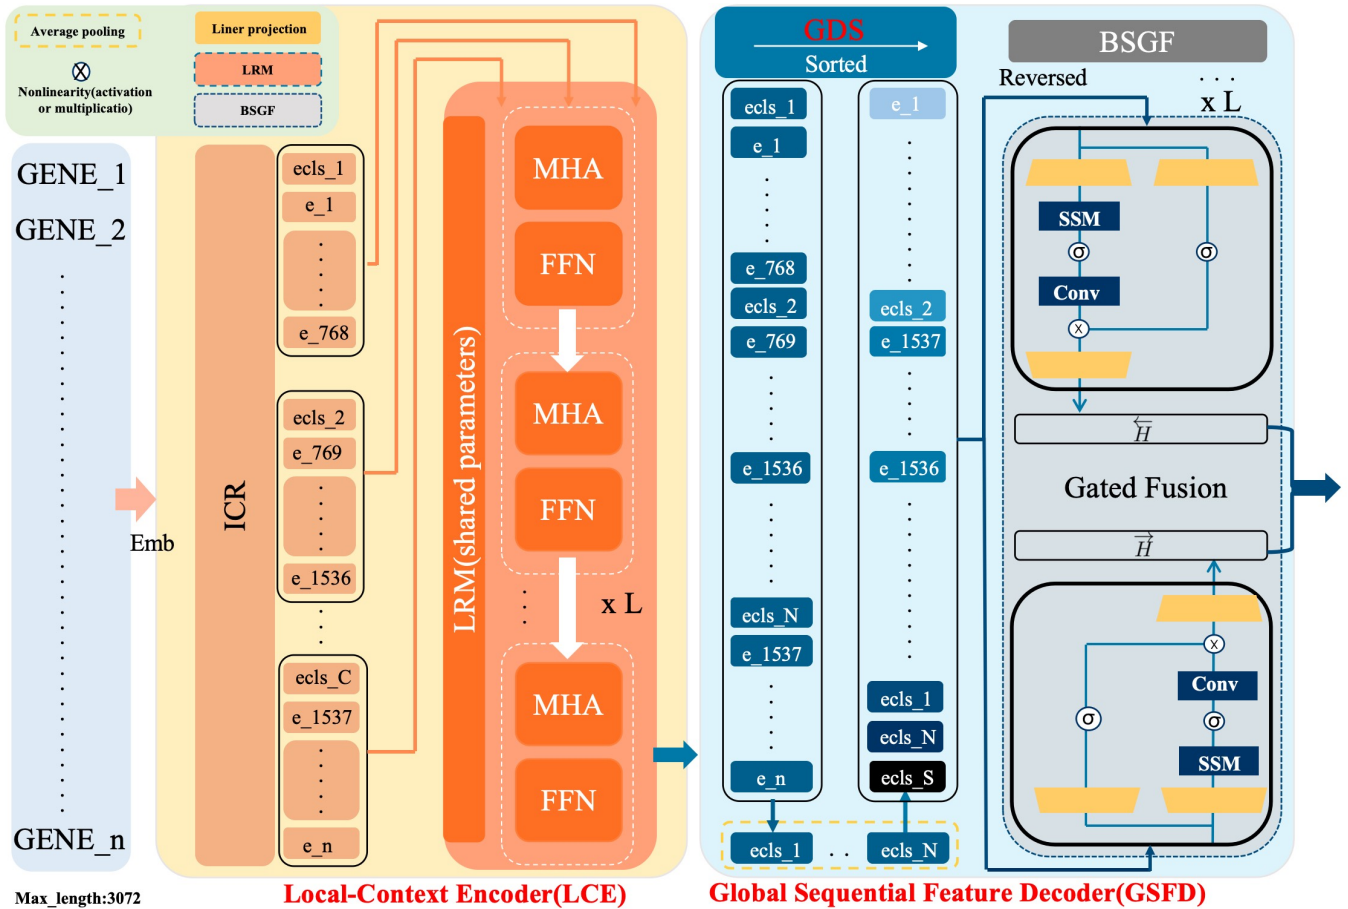


Figure 2: Overview of the PanFoMa algorithm. Firstly, the **Local-context encoder** first divide the input into chunks and encodes local gene interactions within transcriptome chunks using shared Transformer layers. Second, the **global sequential feature decoder** guides dynamic reordering of all genes based on relevance. The reordered sequence is processed by a bidirectional Mamba decoder, and a gating module fuses features to generate the final cell representation.

Module 1: Local-context encoder

This module is designed to efficiently process large-scale gene inputs in parallel and to learn a feature representation for each gene that is rich in information about its local regulatory environment. Its design follows a “**divide and conquer**” strategy, decomposing a large-scale global relationship modeling problem into multiple manageable local problems. It consists of the following two key steps:

Input chunking and representation (ICR) To overcome the quadratic ($O(N^2)$) computational and memory complexity bottleneck faced by standard Transformers when processing ultra-long gene sequences, we first partition the input into chunks. In each training epoch, we randomly sample 3072 genes from a cell and divide them into $C = 4$ non-overlapping chunks, with each chunk k containing a fixed size of $M = 768$ genes. To independently learn and aggregate summary information for each chunk, we prepend an independent, learnable token, $[CLS]$, to the beginning of each chunk.

For each gene $g_{k,i}$ in chunk k , we project its discrete gene

ID and binned expression value through separate embedding layers and then fuse them via element-wise addition to obtain its input feature embedding $e_{k,i} \in \mathbb{R}^D$:

$$e_{k,i} = \text{Emb}_{\text{id}}(g_{\text{id}_{k,i}}) + \text{Emb}_{\text{val}}(g_{\text{val}_{k,i}}) \quad (1)$$

where D is the embedding dimension. Thus, the complete input sequence for the k -th chunk is a matrix $S_k \in \mathbb{R}^{(M+1) \times D}$, represented as $S_k = \{e_{[CLS]_k}, e_{k,1}, \dots, e_{k,M}\}$.

Local relationship modeling (LRM) All C chunks are then processed in parallel by a lightweight Transformer encoder, which consists of a six-layer stack ($L = 6$) with shared parameters across layers. Let the input to layer l for chunk k be $H_k^{(l-1)} \in \mathbb{R}^{(M+1) \times D}$ (where $H_k^{(0)} = S_k$). The output of the layer, $H_k^{(l)}$, is computed by applying the standard Transformer block operations:

$$H_k^{(l)} = \text{TransformerBlock}(H_k^{(l-1)}) \quad (2)$$

After six layers, the module outputs a final hidden state matrix $H_k^{(L)}$ for each chunk. The output of the local-

context encoder is twofold: first, the gene-specific embeddings $H_{\text{genes},k}^{(L)} \in \mathbb{R}^{M \times D}$, which are the rows of $H_k^{(L)}$ corresponding to the gene tokens. Second, the `cls_token` embedding $h_{[\text{CLS}],k}^{(L)} \in \mathbb{R}^D$, which is the first row of $H_k^{(L)}$ and serves as a summary vector for the chunk.

Module 2: Global sequential feature decoder

The core task of this module is to perform global information integration, dynamic sorting, and deep sequential processing based on the locally-informed features generated by the local-context encoder.

Global-informed dynamic sorting (GDS) First, we synthesize a global cell state vector, $h_{\text{global.cls}} \in \mathbb{R}^D$, by applying **Average Pooling** to the `cls_token` summary vectors from all C chunks:

$$h_{\text{global.cls}} = \frac{1}{C} \sum_{k=1}^C h_{[\text{CLS}],k}^{(L)} \quad (3)$$

Next, we concatenate the gene-specific embeddings from all chunks to form a single matrix $H_{\text{genes,full}}^{(L)} \in \mathbb{R}^{N \times D}$, where $N = C \times M$. We then compute an importance score s_i for each gene embedding $h_i \in H_{\text{genes,full}}^{(L)}$ via its dot product with the global cell state vector:

$$s_i = h_i \cdot h_{\text{global.cls}}^T \quad (4)$$

Based on the score vector $s = \{s_1, \dots, s_N\}$, we sort the rows of $H_{\text{genes,full}}^{(L)}$ in descending order to obtain the sorted feature matrix $H_{\text{sorted}}^{(L)}$.

Bidirectional scanning and gated fusion (BSGF) The sorted gene sequence, $H_{\text{sorted}}^{(L)}$, is then fed into a six-layer bidirectional Mamba module. Each layer consists of a forward and a backward Mamba:

$$\begin{aligned} \vec{H}_{\text{mamba}} &= \text{Mamba}_{\text{fwd}}(H_{\text{sorted}}^{(L)}) \\ \overleftarrow{H}_{\text{mamba}} &= \text{Mamba}_{\text{bwd}}(\text{Reverse}(H_{\text{sorted}}^{(L)})) \end{aligned}$$

where $\vec{H}_{\text{mamba}}, \overleftarrow{H}_{\text{mamba}} \in \mathbb{R}^{N \times D}$ are the output hidden state sequences.

To fuse these bidirectional flows, a **gated fusion mechanism** learns a gate vector $\gamma_i \in \mathbb{R}^D$ for each gene’s sorted feature $h_{\text{sorted},i} \in H_{\text{sorted}}^{(L)}$:

$$\gamma_i = \sigma(\text{Linear}(h_{\text{sorted},i})) \quad (5)$$

where σ is the Sigmoid function. This gate is used to compute the final fused feature $h_{\text{fused},i}$ as a weighted sum of the forward and backward hidden states, $\vec{h}_{\text{mamba},i}$ and $\overleftarrow{h}_{\text{mamba},i}$:

$$h_{\text{fused},i} = \gamma_i \odot \vec{h}_{\text{mamba},i} + (1 - \gamma_i) \odot \overleftarrow{h}_{\text{mamba},i} \quad (6)$$

where \odot denotes element-wise multiplication. This approach allows the model to adaptively weigh the contextual information from both directions for each gene.

Pan-Cancer benchmark construction

To construct a diverse and high-quality Pan-Cancer benchmark, named as *PanFoMaBench*, we followed a rigorous and systematic workflow. We began by extensively searching and curating publicly available human cancer single-cell transcriptomic datasets from the National Center for Biotechnology Information (NCBI) database. Selected datasets were required to include essential metadata, such as cellular transcriptomic profiles, cancer type, tissue of origin, and corresponding patient information. To ensure consistency and enable cross-dataset integration, all gene identifiers were standardized to official symbols defined by the HUGO Gene Nomenclature Committee (HGNC).

After standardization, the individual datasets were merged to form a unified Pan-Cancer cohort. This integrated dataset was then subjected to a *comprehensive quality control (QC) pipeline to ensure data fidelity and reliability*. The QC process involved the following filtering steps:

- (1) Removal of cells with a low number of expressed genes;
- (2) Exclusion of potential doublets or multiplets, identified by abnormally high gene or UMI counts;
- (3) Filtering out cells with a high proportion of mitochondrial gene expression, indicative of low viability;
- (4) Elimination of genes with low expression across the dataset to reduce noise and improve computational efficiency.

To support the training and validation of our model, and to contribute a valuable resource to the pan-cancer research community, we constructed a large-scale, high-quality pan-cancer single-cell transcriptomic dataset. We began by conducting a systematic literature search in public databases, including the National Center for Biotechnology Information (NCBI), using keywords such as "pan-cancer" and "cancer". This search process ultimately identified 83 relevant published studies, from which the publicly available single-cell sequencing data were sourced for our dataset. Following rigorous data filtering, standardized preprocessing pipelines, and uniform quality control, we successfully integrated and constructed this high-quality benchmark dataset, which was used for all subsequent analyses. In total, our curated dataset comprises approximately 3.5 million high-quality cells from 616 patients, comprehensively covering 34 distinct cancer subtypes and originating from 23 different tissue types.

Experiments

Experiment setting

This section aims to comprehensively assess the proposed PanFoMa from two perspectives: Pan-cancer diagnosis and foundation model capabilities. Specifically, Pan-cancer diagnosis is evaluated via the task of Pan-cancer classification, while the foundation model capabilities are assessed through Batch integration, cell type annotation, and multi-omic integration.

We compare our methods with the recent proposed single-cell foundation models to demonstrate our better cell representation learning ability: GeneFormer (Cui et al. 2024b),

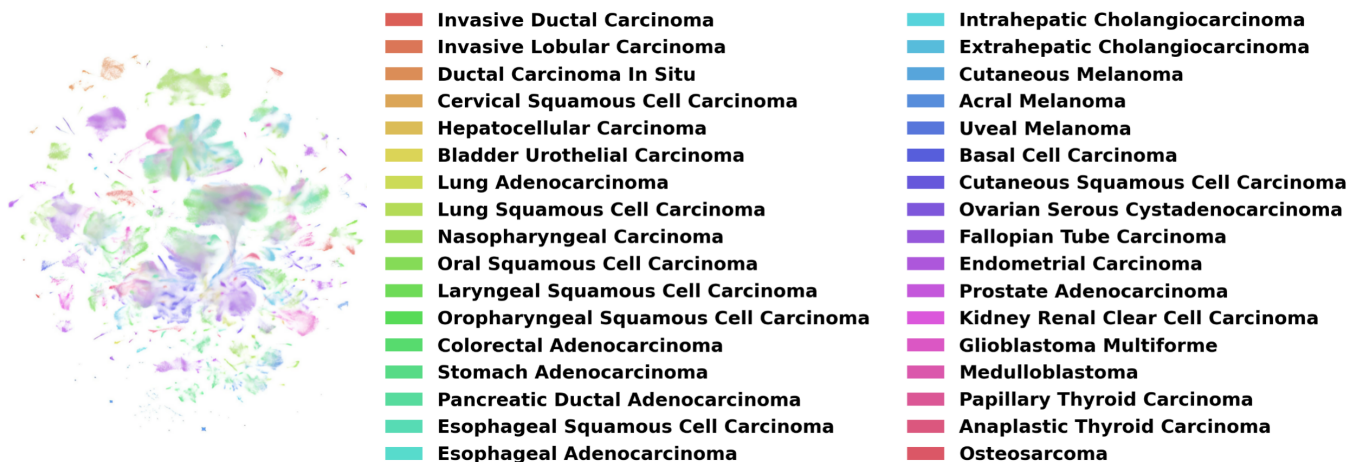


Figure 3: Visualization of the cancer subtype distribution in PanFoMaBench.

scGPT (Cui et al. 2024a), scFoundation (Hao et al. 2024) and GeneMamba (Qi et al. 2025).

Pan-cancer diagnosis

Pan-Cancer diagnosis is a cancer subtype classification problem (Ocasio et al. 2024). To assess the performance of our proposed PanFoMa on the Pan-Cancer diagnosis task, we fine-tuned the model on the training datasets of our benchmark and evaluated its classification accuracy on the corresponding test sets. For comparison, four state-of-the-art single-cell foundation models were evaluated under identical settings.

Datasets	Models	ACC	Macro-F1
PanFoMaBench	GeneFormer	0.9124	0.8851
	scGPT	0.9013	0.8732
	scFoundation	0.8876	0.8491
	GeneMamba	0.9026	0.8619
	Ours	0.9474	0.9250

Table 1: Pan-Cancer diagnosis on our benchmark.

As detailed in Table 1, PanFoMa achieves a remarkable accuracy of 94.74%, surpassing the second-best model, GeneFormer, by a margin of 3.5%. Regarding the Macro-F1 score, PanFoMa attains 92.5%, marking a substantial 4.0% improvement over GeneFormer. Notably, our model consistently outperforms GeneMamba across both metrics. These findings underscore the critical importance of integrating both global and local contextual information for precise Pan-Cancer diagnosis. The inherent complexity of this task is visualized in Figure 3, which illustrates the diverse distribution of the 34 cancer subtypes within PanFoMaBench.

Batch integration

To evaluate batch integration performance, we follow the protocol in (Qi et al. 2025), adopting two widely used metrics: *Avg_batch* (assessing batch effect removal)

and *Avg_bio* (measuring biological variance preservation). Higher scores indicate better performance. We benchmark PanFoMa against four advanced models on five challenging, multi-batch public datasets. As summarized in Table 2, PanFoMa consistently outperforms both transformer-based and Mamba-based baselines in most cases, demonstrating its superior capability in integrating multi-batch single-cell data.

Cell type annotation

Cell type annotation (Chen et al. 2023) aims to assign a known biological identity label to each cell based on its unique gene expression profile. To comprehensively evaluate the performance of our proposed PanFoMa on the cell type annotation task, we selected four widely used public benchmark datasets (Qi et al. 2025; Hoadley et al. 2018)—hPancreas, MS, Myeloid, and Myeloid_b—and adopted Accuracy and Macro-F1 score as the primary evaluation metrics.

We conducted a direct comparison between PanFoMa and several state-of-the-art baseline models. As summarized in Table 3, PanFoMa outperforms scGPT by 1.0% on hPancreas, 0.9% on MS, 1.7% on Myeloid, and 3.0% on Myeloid_b. Compared to GeneMamba, PanFoMa achieves improvements of 1.0% on hPancreas, 7.4% on MS, and 1.2% on Myeloid_b. These consistent gains in accuracy demonstrate the effectiveness of PanFoMa in leveraging both global and local contextual information for more accurate and robust cell type annotation.

Multi-omic integration

The multi-omic integration task (Cui et al. 2024a) aims to unify data from different molecular modalities (e.g., transcriptome, epigenome, proteome) into a shared embedding space that reflects true biological states rather than technical variation. To evaluate the integration capability of PanFoMa, we conducted experiments on three representative public datasets (He et al. 2024; Cui et al. 2024a): 10x Multi-ome PBMC (RNA + ATAC), BMBC (RNA + protein), and

Metric	Datasets	GeneFormer	scGPT	scFoundation	GeneMamba	Ours
Avg_batch	Immune	0.8153	0.9194	0.8904	0.9536	0.9641
	PBMC12k	0.9545	0.9755	0.9628	0.9604	0.9701
	BMMC	0.7720	0.8431	0.7598	0.9157	0.9312
	Perirhinal Cortex	0.9127	0.9600	0.9560	0.9673	0.9661
	Covid-19	0.8240	0.8625	0.8346	0.8742	0.9173
Avg_bio	Immune	0.6983	0.7879	0.7337	0.8131	0.8332
	PBMC12k	0.7891	0.9018	0.8662	0.8344	0.8631
	BMMC	0.6324	0.6576	0.5250	0.7628	0.8021
	Perirhinal Cortex	0.8547	0.9552	0.9606	0.9062	0.9598
	Covid-19	0.5567	0.6476	0.5468	0.5537	0.6338

Table 2: Comparison of batch integration across datasets.

Datasets	Models	Acc	Macro-F1
hPancreas	GeneFormer	0.9665	0.7450
	scGPT	0.9710	0.7632
	scFoundation	0.9602	0.7101
	GeneMamba	0.9713	0.7710
	Ours	0.9815	0.7760
MS	GeneFormer	0.7650	0.6220
	scGPT	0.8471	0.6630
	scFoundation	0.7763	0.6812
	GeneMamba	0.6825	0.5342
	Ours	0.8563	0.7016
Myeloid	GeneFormer	0.6445	0.3600
	scGPT	0.6341	0.3562
	scFoundation	0.6446	0.3646
	GeneMamba	0.6607	0.3650
	Ours	0.6515	0.3529
Myeloid.b	GeneFormer	0.9540	0.9380
	scGPT	0.9421	0.9434
	scFoundation	0.9574	0.9569
	GeneMamba	0.9603	0.9235
	Ours	0.9726	0.9351

Table 3: The cell type annotation accuracy across datasets.

the more challenging mosaic ASAP PBMC dataset.

We adopted the average biological variance preservation score (Avg_bio) as the main evaluation metric, where higher values indicate better retention of biological signal post-integration. We compared PanFoMa against leading methods, including scGPT (Cui et al. 2024a), scGLUE (Cao and Gao 2022), Seurat v4 (Hao et al. 2021), and scMoMaT (Zhang et al. 2023).

As shown in Table 4, PanFoMa achieves state-of-the-art performance, with gains of 3.1% on Multiome PBMC and 2.4% on BMMC over the best baselines. These results highlight PanFoMa’s ability to capture both local and global biological context, enabling more accurate and biologically meaningful multi-omic integration.

Datasets	Models	Avg_bio
10x Multiome PBMC	scGPT	0.758
	scGLUE	0.747
	Seurat v4	0.722
	scMoMaT	0.725
	Ours	0.789
BMMC (RNA + Protein)	scGPT	0.697
	scGLUE	0.600
	Seurat v4	0.650
	scMoMaT	0.630
	Ours	0.721
ASAP PBMC	scGPT	0.587
	scGLUE	0.561
	Seurat v4	0.541
	scMoMaT	0.546
	Ours	0.579

Table 4: Comparison of multi-omic integration across datasets.

Conclusion

In this study, we introduce PanFoMa, a lightweight hybrid neural network that effectively balances performance and efficiency for pan-cancer single-cell representation learning. Our model introduces a sophisticated hierarchical "local-to-global" design paradigm. This design successfully decouples the complex task of transcriptome modeling into two independent sub-tasks: parallel local context encoding and efficient global information integration. To rigorously evaluate our model, we constructed a large-scale pan-cancer single-cell benchmark, PanFoMaBench, spanning 34 cancer subtypes and over 3.5 million high-quality cells. Extensive experiments demonstrate that PanFoMa outperforms existing state-of-the-art models across a wide range of tasks, including cell-type annotation, batch correction and multi-omic integration. These results highlight the potential of PanFoMa as a foundational model for advancing precision oncology and single-cell computational biology.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 62271237), the Key Project of the Natural Science Foundation of Jiangxi Province (Grant No. 20242BAB26014), the 2025 Pilot Project of Jiangxi Future Health Interdisciplinary Center (Grant No. 9167-28280006), the Hunan Natural Science Foundation (Grant No. 2025JJ50338), and the Shanghai Education Committee AI Project (Grant No. JWAIYB-2).

References

- Adduri, A.; Gautam, D.; Bevilacqua, B.; Imran, A.; Shah, R.; Naghipourfar, M.; Teyssier, N.; Ilango, R.; Nagaraj, S.; Ricci-Tam, C.; et al. 2025. Predicting cellular responses to perturbation across diverse contexts with STATE. *bioRxiv*, 2025–06.
- Bjerregaard-Michelsen, S.; Poulsen, L.; Bjerrum, A.; Bøgsted, M.; and Vesteghem, C. 2025. Machine learning for prediction of 30-day mortality in patients with advanced cancer: comparing pan-cancer and single-cancer models. *ESMO Real World Data and Digital Oncology*, 8: 100146.
- Cao, Z.-J.; and Gao, G. 2022. Multi-omics single-cell data integration and regulatory inference with graph-linked embedding. *Nature Biotechnology*, 40(10): 1458–1466.
- Chen, J.; Xu, H.; Tao, W.; Chen, Z.; Zhao, Y.; and Han, J.-D. J. 2023. Transformer for one stop interpretable cell type annotation. *Nature Communications*, 14(1): 223.
- Cui, H.; Wang, C.; Maan, H.; Pang, K.; Luo, F.; Duan, N.; and Wang, B. 2024a. scGPT: toward building a foundation model for single-cell multi-omics using generative AI. *Nature methods*, 21(8): 1470–1480.
- Cui, Z.; Xu, T.; Wang, J.; Liao, Y.; and Wang, Y. 2024b. Geneformer: Learned gene compression using transformer-based context modeling. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8035–8039. IEEE.
- Dutta, A. K.; Alberge, J.-B.; Sklaventis-Pistofidis, R.; Lightbody, E. D.; Getz, G.; and Ghobrial, I. M. 2022. Single-cell profiling of tumour evolution in multiple myeloma—opportunities for precision medicine. *Nature Reviews Clinical Oncology*, 19(4): 223–236.
- Fang, C.; Cui, W.; Hu, Z.; Liu, W.; Chen, S.; Chang, S.; Long, Q.; Li, C.; Liu, Y.; Jiang, H.; et al. 2025. Cell-GraphCompass: Modeling Single Cells with Graph Structure Foundation Model. *National Science Review*, nwaf255.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Hao, M.; Gong, J.; Zeng, X.; Liu, C.; Guo, Y.; Cheng, X.; Wang, T.; Ma, J.; Zhang, X.; and Song, L. 2024. Large-scale foundation model on single-cell transcriptomics. *Nature methods*, 21(8): 1481–1491.
- Hao, Y.; Hao, S.; Andersen-Nissen, E.; Mauck, W. M.; Zheng, S.; Butler, A.; Lee, M. J.; Wilk, A. J.; Darby, C.; Zager, M.; et al. 2021. Integrated analysis of multimodal single-cell data. *Cell*, 184(13): 3573–3587.
- He, Z.; Dony, L.; Fleck, J. S.; Szałata, A.; Li, K. X.; Slišković, I.; Lin, H.-C.; Santel, M.; Atamian, A.; Quadrato, G.; et al. 2024. An integrated transcriptomic cell atlas of human neural organoids. *Nature*, 635(8039): 690–698.
- Hoadley, K. A.; Yau, C.; Hinoue, T.; Wolf, D. M.; Lazar, A. J.; Drill, E.; Shen, R.; Taylor, A. M.; Cherniack, A. D.; Thorsson, V.; et al. 2018. Cell-of-origin patterns dominate the molecular classification of 10,000 tumors from 33 types of cancer. *Cell*, 173(2): 291–304.
- Huang, X.; Huang, Z.; Li, S.; Qu, W.; He, T.; Hou, Y.; Zuo, Y.; and Ouyang, W. 2024. Frozen clip transformer is an efficient point cloud encoder. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2382–2390.
- Huang, X.; Huang, Z.; Zuo, Y.; Gong, Y.; Zhang, C.; Liu, D.; and Fang, Y. 2025. PSReg: Prior-guided Sparse Mixture of Experts for Point Cloud Registration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 3788–3796.
- Jovic, D.; Liang, X.; Zeng, H.; Lin, L.; Xu, F.; and Luo, Y. 2022. Single-cell RNA sequencing technologies and applications: A brief overview. *Clinical and translational medicine*, 12(3): e694.
- Lareau, C. A.; Parker, K. R.; and Satpathy, A. T. 2021. Charting the tumor antigen maps drawn by single-cell genomics. *Cancer cell*, 39(12): 1553–1557.
- Mei, G.; Huang, X.; Liu, J.; Zhang, J.; and Wu, Q. 2022. Unsupervised point cloud pre-training via contrasting and clustering. In *2022 IEEE International Conference on Image Processing (ICIP)*, 66–70. IEEE.
- Nofech-Mozes, I.; Soave, D.; Awadalla, P.; and Abelson, S. 2023. Pan-cancer classification of single cells in the tumour microenvironment. *Nature Communications*, 14(1): 1615.
- Nopour, R. 2024. Prediction of 12-month recurrence of pancreatic cancer using machine learning and prognostic factors. *BMC Medical Informatics and Decision Making*, 24(1): 339.
- Ocasio, B. A.; Hu, J.; Stathias, V.; Martinez, M. J.; Burnstein, K. L.; and Schürer, S. C. 2024. Pan-Cancer Drug Sensitivity Prediction from Gene Expression using Deep Learning. *bioRxiv*.
- Ouyang, S.; Zhang, J.; Liu, F.; Jiang, Q.; Xing, W.; Chen, J.; and Zhang, J. 2025. Global research trends and hotspots in prognostic prediction models for pancreatic cancer: a bibliometric analysis. *Frontiers in Oncology*, 15: 1588735.
- Qi, C.; Fang, H.; Hu, T.; Jiang, S.; and Zhi, W. 2025. Bidirectional Mamba for Single-Cell Data: Efficient Context Learning with Biological Fidelity. *arXiv preprint arXiv:2504.16956*.
- Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I.; et al. 2018. Improving language understanding by generative pre-training.
- Salmanpour, M. R.; Mehrnia, S. S.; Ghandilu, S. J.; Safahi, Z.; Falahati, S.; Taeb, S.; Mousavi, G.; Maghsoudi, M.; Shariftabrizi, A.; Hacihaliloglu, I.; and Rahmim, A. 2025. Handcrafted vs. Deep Radiomics vs. Fusion vs. Deep Learning: A Comprehensive Review of Machine Learning -Based Cancer Outcome Prediction in PET and SPECT Imaging.

Singh, D. P.; Kour, P.; Banerjee, T.; and Swain, D. 2025. A Comprehensive Review of Various Machine Learning and Deep Learning Models for Anti-Cancer Drug Response Prediction: Comparative Analysis With Existing State of the Art Methods. *Archives of Computational Methods in Engineering*, 1–25.

Štancl, P.; and Karlić, R. 2023. Machine learning for pan-cancer classification based on RNA sequencing data. *Frontiers in molecular biosciences*, 10: 1285795.

Theus, A.; Barkmann, F.; Wissel, D.; and Boeva, V. 2024. CancerFoundation: A single-cell RNA sequencing foundation model to decipher drug resistance in cancer. In *NeurIPS 2024 Workshop on AI for New Drug Modalities*.

Van de Sande, B.; Lee, J. S.; Mutasa-Gottgens, E.; Naughton, B.; Bacon, W.; Manning, J.; Wang, Y.; Pollard, J.; Mendez, M.; Hill, J.; et al. 2023. Applications of single-cell RNA sequencing in drug discovery and development. *Nature Reviews Drug Discovery*, 22(6): 496–520.

Weinstein, J. N.; Collisson, E. A.; Mills, G. B.; Shaw, K. R.; Ozenberger, B. A.; Ellrott, K.; Shmulevich, I.; Sander, C.; and Stuart, J. M. 2013. The cancer genome atlas pan-cancer analysis project. *Nature genetics*, 45(10): 1113–1120.

Yang, P.; Yin, J.; Zhang, G.; Li, X.; Chen, T.; Zhao, W.; Tang, J.; Lv, L.; and Lv, X. 2025. Comprehensive pan-cancer analysis identified SLC16A3 as a potential prognostic and diagnostic biomarker. *Cancer Cell International*, 25(1): 168.

Zeng, Y.; Xie, J.; Shangguan, N.; Wei, Z.; Li, W.; Su, Y.; Yang, S.; Zhang, C.; Zhang, J.; Fang, N.; et al. 2025. CellFM: a large-scale foundation model pre-trained on transcriptomics of 100 million human cells. *Nature Communications*, 16(1): 4679.

Zhang, B.; Li, S.; Jian, J.; Ren, X.; Zhao, Z.; Guo, L.; Su, F.; Meng, Z.; and Zhao, Z. 2025. From Single-Cancer to Pan-cancer Prognosis: A Multi-modal Deep Learning Framework for Survival Analysis with Robust Generalization Capability. *The American Journal of Pathology*.

Zhang, Z.; Sun, H.; Mariappan, R.; Chen, X.; Chen, X.; Jain, M. S.; Efremova, M.; Teichmann, S. A.; Rajan, V.; and Zhang, X. 2023. scMoMaT jointly performs single cell mosaic integration and multi-modal bio-marker detection. *Nature Communications*, 14(1): 384.

Zheng, X.; Huang, X.; Mei, G.; Hou, Y.; Lyu, Z.; Dai, B.; Ouyang, W.; and Gong, Y. 2024. Point cloud pre-training with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22935–22945.