Towards Self-Supervised Foundation Models for Critical Care Time Series

Anonymous Author(s)
Affiliation
Address
email

Abstract

Domain-specific foundation models for healthcare have expanded rapidly in recent years, yet foundation models for critical care time series remain relatively underexplored due to the limited size and availability of datasets [1]. In this work, we introduce an early-stage pre-trained foundation model for critical care time-series based on the Bi-Axial Transformer (BAT) [2], trained on pooled electronic health record datasets. We demonstrate effective transfer learning by fine-tuning the model on a dataset distinct from the training sources for mortality prediction, where it outperforms supervised baselines, particularly in low-data regimes ($\sim 10,000$ samples, and most notably $< 5,000$). These contributions highlight the potential of self-supervised foundation models for critical care times series to support generalizable and robust clinical applications in resource-limited settings.

1 Introduction and related works

Foundation models built on Transformers [3] have achieved remarkable success across several domains, including natural language processing [4, 5] and computer vision [6, 7, 8]. Despite this success, general-purpose foundation models often perform poorly on healthcare applications, which are impeded by the complexities of medical data and the scarcity of publicly available labeled datasets [1]. Consequently, the development of healthcare-specific foundation models has accelerated rapidly since 2018, targeting diverse applications such as clinical natural language processing [9, 10], medical imaging [11, 12], omics analysis [13, 14], video and audio interpretation [15, 16] and multi-modality [17, 18]. Foundation models for Electronic Health Records (EHRs) have primarily focused on structured EHR data, such as modelling and predicting ICD-10 codes [19, 20]. Although these models can scale to large patient populations, they remain limited in their ability to capture physiological patterns [21]. One area that remains relatively underexplored is the development of foundation models for critical care time-series data [22]. Models trained in this domain are often hindered by issues of reproducibility [23], limited by simple learning paradigms that rely on one supervised task [24, 25, 26] or are trained on a small, homogeneous dataset [27, 28, 29]. It has been demonstrated that these models do not transfer well to new clinical settings [22]. Recent work from Burger et al. [21] are the first to combine pooling of several critical care time-series datasets with self-supervised pretraining, and thereby delivering a critical care time series foundation model, ICareFM [21]. Both the model and code are currently not open-source.

To push critical care time-series data towards foundation model development, we modify the Bi-axial Transformer (BAT) architecture, presented in the work of DeVries et al. [2], for self-supervised pretraining and conduct all experiments within the Yet Another ICU Benchmark (YAIB) framework [23], ensuring transparency, reproducibility and pooling of datasets. Our contributions are as follows: 1) We release the first ICU-specific model for foundational capabilities with an open-source, reproducible repository, 2) we demonstrate its ability to transfer effectively to an unseen dataset distinct from the training sources and a new downstream clinical task, outperforming supervised baselines, and 3) we

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

show that it performs especially well in low-data regimes, highlighting the potential of self-supervised pre-training for resource-limited clinical settings that may not achieve robust performance from supervised models.

2 Methodology

Let $\mathcal{D} := \bigcup_{k=1}^{K} \mathcal{D}_k$, where a critical care dataset $\mathcal{D}_k$ is defined as a set of tuples, $\{(\mathbf{X}_i,\ \mathbf{p}_i,\ \mathbf{h}_i)\}_{i=1}^{N_k}$. Every patient $i \in N_k$ is represented with a multivariate time series of measurements $\mathbf{X}_i \in \mathbb{R}^{T_i \times D}$ from monitoring devices and laboratory tests, with K datasets containing the same measurement types D. As each $\mathbf{X}_i$ is irregularly sampled across both t and d, we also record measurement times $\mathbf{h}_i \in \mathbb{R}^{T_i}$. The time series is supplemented with patient demographics $\mathbf{p}_i$ which are fixed for the duration of T.

In critical care applications, it is often of interest to predict a class label, such as $y_i \in \{0, 1\}$ for mortality status, given $(\mathbf{X}_i, \mathbf{p}_i, \mathbf{h}_i)$. Popular approaches split a single dataset $\mathcal{D}_1$ into train/test/validation subsets, and perform supervised training to predict y_i. We hypothesize that combining and pretraining on additional datasets $\{\mathcal{D}_2, ...\mathcal{D}_K\}$, in a self-supervised fashion will result in a model that learns richer, more generalized representations of sensor identities and their measurement distributions.

2.1 Pretraining and Fine-tuning

We begin with self-supervised pretraining with forecasting as the prediction task. Specifically, $\mathbf{X}_i$ sampled from the auxiliary datasets $\{\mathcal{D}_2, ...\mathcal{D}_K\}$ are split into an observation window $\mathbf{X}_i^{\text{obs}} \in \mathbb{R}^{T^{\text{obs}} \times D}$ and a forecasting window $\mathbf{X}_i^{\text{for}} \in \mathbb{R}^{T^{\text{for}} \times D}$. The learning objective is to predict $\mathbf{X}_i^{\text{for}}$ given $(\mathbf{X}_i^{\text{obs}}, \mathbf{p}_i, \mathbf{h}_i^{\text{obs}})$. Due to the irregular sampling across h_i and sparsity in X_i, the forecasting must be sample-specific. To handle sparsity and ensure temporal alignment, we leverage a masked loss:

$$\mathcal{L}^{\text{Pre}} = \frac{1}{\sum_{k=2}^{K} N_k} \sum_{k=2}^{K} \sum_{i=1}^{N_k} \left\| \mathbf{M}_i^{\text{for}} \odot \left(\hat{\mathbf{X}}_i^{\text{for}} - \mathbf{X}_i^{\text{for}} \right) \right\|_F^2$$

where the mask $\mathbf{M}_i^{\text{for}}$ ensures that only observed values in the forecasting window contribute to the loss.

We then fine-tune and test on $\mathcal{D}_1$ with supervised learning, where the objective is our initial goal of predicting the mortality status $y_i \in \{0, 1\}$ as a binary classification task. We used the binary cross-entropy loss:

$$\mathcal{L}^{\text{Fine}} = -\frac{1}{N_1} \sum_{i=1}^{N_1} \Big[y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i) \Big]$$

2.2 Bi-axial Transformer (BAT)

The Bi-Axial Transformer (BAT) [2] is a generalization of the Axial Transformer [30] with parallel axial attention along each dimension. BAT attends to both the temporal and clinical feature axes through axial attention mechanisms, and explicitly accounts for missing values. The model architecture is described in detail in Appendix A.1. BAT shows promising results in prediction tasks for clinical irregular multivariate time series data with state-of-the-art performance for sepsis classification, and competitive performance on mortality prediction when compared to several models. Investigation into model attention maps revealed evidence of BAT learning from informative missingness, and it showed an increased robustness to sparsity in comparison to other Transformer models. These properties make BAT a strong candidate for modeling clinical multivariate time series. BAT's prediction head was adapted in this work to support both binary classification, outputting $\hat{y} \in \{0, 1\}$, and forecasting, outputting $\hat{\mathbf{X}}^{\text{for}} \in \mathbb{R}^{T^{\text{for}} \times D}$, as illustrated in Figure 2.

2.3 Yet another ICU benchmark (YAIB)

Yet another ICU benchmark (YAIB) provided by van de Water et al. [23] is a modular, end-to-end framework supporting transparent benchmarking for clinical machine learning on ICU data (Appendix A.2). YAIB supports several publicly available EHR datasets and enables cohort selection, harmonization, and a variety of supervised learning tasks, with several implemented machine learning and deep learning models. This work has extended the YAIB framework to contain a self-supervised learning module, detailed in Appendix A.3, as well as the previously described BAT architecture. Building on YAIB's emphasis on transparent and reproducible benchmarking, we use the framework to evaluate our self-supervised models in a controlled setting.

3 Experiments

All experiments were performed using the Medical Information Mart for Intensive Care Database versions III and IV (MIMIC-III & IV, [31, 32]) and The eICU Collaborative Research Database (eICU [33]). Further information about the three datasets and their preprocessing can be found in Appendix A.4. A t-SNE [34] analysis of the preprocessed and harmonized datasets was performed, and Figure 4 indicates that they occupy a similar overall distribution. Despite the overlap in data distribution, we reproduce similar findings to previous studies [22, 35] showing that training models on one dataset does not transfer effectively to new ones during inference (Figure 5). Hyperparameters for all models can be found in Appendix A.6.

3.1 Baseline Comparison

BAT was pre-trained with self-supervised forecasting on a pooled version of the datasets, with one held-out for fine-tuning. Pre-training followed the cross-validation setup described in Appendix A.7, and details on the models can be found in Appendix A.8. The pre-trained models were evaluated on the ability to be fine-tuned and generalize to the held-out dataset for mortality prediction. To test model robustness, fine-tuning was performed on varying subsets of the held-out data, with training set sizes ranging from 100 to 9,506 samples.

Figure 1 illustrates the results of pre-training on the pooled dataset consisting of eICU & MIMIC-IV and fine-tuning on MIMIC-III. The mean and standard deviations of AUC-PR and AUC-ROC over all training set sizes reflects five random subsamples of the training sets with preserved class imbalance. The pre-trained model was compared to two baseline models trained from scratch in a supervised manner: one using the same architecture (BAT) and another based on a vanilla Transformer that does not natively handle missing values and instead relies on mean-imputed data. The AUC-PR performances, and partially the AUC-ROC performance, revealed that the pre-trained model outperformed the baseline models on all training set sizes > 500 samples, but most significantly on dataset sizes < 5000 samples. Furthermore, fine-tuning only the binary classification head of the pre-trained model yields close to similar performance when compared to full-model fine-tuning. Similar results were seen when MIMIC-IV and eICU were held out (Table 1). In these two cases, the BAT model trained from scratch outperformed the fine-tuned version where only the classification head was fine-tuned, when enough labeled data was available (≥ 3000 and ≥ 5000 samples, respectively). This suggests that pre-training on pooled eICU+MIMIC-IV data leads to more robust representations compared to the other pooled pre-training datasets. Results for held-out test sets of additional sizes are provided in Appendix A.9, and show similar performance.

4 Discussion

We evaluate BAT as a model architecture for a critical care time-series foundation model and find that it outperforms initial supervised baselines on held-out datasets distinct from the training sources for mortality prediction, with training set sizes up to 10,000 samples. Fine-tuning only the binary classification head of the pre-trained model achieves performance comparable to full-model fine-tuning, suggesting that the learned embeddings are both informative and transferable to downstream tasks on unseen datasets. This effect was most pronounced for the model pretrained on the largest dataset (277K samples; see Appendix A.8), MIMIC-IV + eICU, and fine-tuned on MIMIC-III. By contrast, the other two models with smaller training set sizes were outperformed by

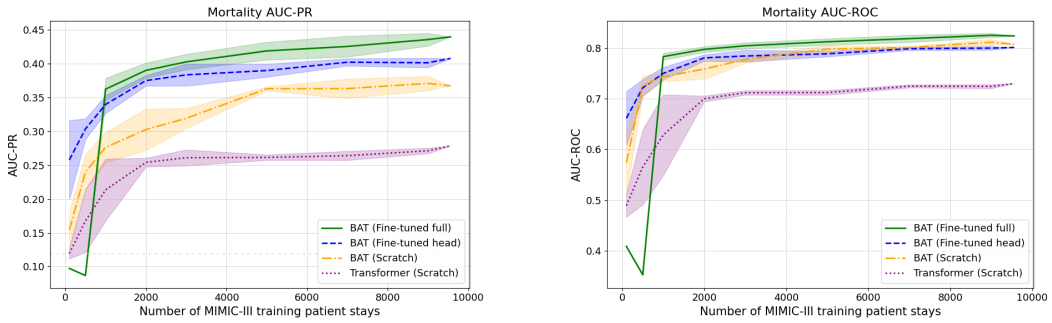

(a) **(b)**

Figure 1: Performance of fine-tuned model (pre-trined on eICU + MIMIC-IV) and supervised models trained from scratch on MIMIC-III. **(a)** AUC-PR (grey dashed line indicates the positive class prevalence) and **(b)** AUC-ROC across training set sizes ranging from 100 to 9,506 samples. Models include BAT full & head fine-tuning, and models trained from scratch: BAT and Transformer.

Table 1: Average model performance, (AUC-PR $\pm$ sd) across multiple dataset sizes from MIMIC-III, MIMIC-IV, and eICU. The pre-trained BAT models are fine-tuned and the baseline models are trained from scratch on the subsets. Highest performance for each dataset size is in bold, second highest is underlined.

Fine-tuning/ Training dataset	Dataset size	BAT (Fine-tuned full)	BAT (Fine-tuned head)	BAT (Scratch)	Transformer (Scratch)
MIMIC-III	1000	**36.24 $\pm$ 1.63**	33.98 $\pm$ 1.32	27.63 $\pm$ 2.23	21.34 $\pm$ 4.58
	5000	**41.89 $\pm$ 1.31**	38.99 $\pm$ 0.96	36.30 $\pm$ 0.31	26.14 $\pm$ 0.40
	9000	**43.57 $\pm$ 0.94**	40.14 $\pm$ 0.70	37.09 $\pm$ 1.06	27.13 $\pm$ 0.41
MIMIC-IV	1000	**28.98 $\pm$ 0.85**	26.97 $\pm$ 1.71	26.12 $\pm$ 1.95	13.06 $\pm$ 1.56
	5000	**38.10 $\pm$ 1.36**	31.31 $\pm$ 1.17	34.97 $\pm$ 1.03	18.00 $\pm$ 1.11
	9000	**38.75 $\pm$ 0.53**	32.19 $\pm$ 1.00	38.41 $\pm$ 0.88	18.91 $\pm$ 1.36
eICU	1000	**28.37 $\pm$ 1.13**	25.39 $\pm$ 1.56	20.86 $\pm$ 3.31	6.58 $\pm$ 4.00
	5000	**33.89 $\pm$ 0.49**	29.09 $\pm$ 0.62	30.13 $\pm$ 1.37	14.41 $\pm$ 0.42
	9000	**35.20 $\pm$ 0.87**	29.53 $\pm$ 0.95	31.59 $\pm$ 0.94	14.61 $\pm$ 1.06

the equivalent baseline models trained from scratch when only the classification head was fine-tuned and sufficient labeled data was available. These findings indicate that large and diverse pre-training datasets are crucial for learning representations that generalize and transfer effectively across datasets. Together, these results highlight the value of such a model in clinical settings where labeled data and computational resources are limited. Overall, this work demonstrates the feasibility and benefits of training foundation models for critical care time-series data within a transparent, reproducible framework [23], extended for self-supervised pre-training to enable transfer learning and more robust, generalizable models.

4.1 Limitations

Our experiments were limited to MIMIC-III, MIMIC-IV, and eICU. To reach a training set size that is competitive with other foundation models, future work will need to incorporate additional datasets, such as those presented in Burger et al. [22]. Compared with fields like natural language processing, the number and size of publicly available ICU and critical care datasets remain very limited. Expanding the training data may therefore also require incorporating time-series datasets from other domains, such as weather [36] or electricity consumption [37]. This would allow us to assess whether exposure to broader time-series distributions improves model performance, or if domain-specific data is necessary given the sparse and irregular nature of critical care records, as suggested by other healthcare foundation models [1].

References

[1] Wasif Khan, Seowung Leem, Kyle B See, Joshua K Wong, Shaoting Zhang, and Ruogu Fang. A comprehensive survey of foundation models in medicine. *Ieee Reviews in Biomedical Engineering*, PP(99):1–22, 2025. ISSN 19411189, 19373333. doi: 10.1109/RBME.2025.3531360.

[2] Rachael DeVries, Casper Christensen, Marie Lisandra Zepeda Mendoza, and Ole Winther. Bi-axial transformers: Addressing the increasing complexity of ehr classification, 2025. URL https://arxiv.org/abs/2508.12418.

[3] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019. URL https://arxiv.org/abs/1810.04805.

[5] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. *OpenAI Technical Report*, 2018. URL https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf.

[6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. URL https://arxiv.org/abs/2010.11929.

[7] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows, 2021. URL https://arxiv.org/abs/2103.14030.

[8] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers distillation through attention, 2021. URL https://arxiv.org/abs/2012.12877.

[9] Chen Zeming, Romanou Angelika, Bonnet Antoine, Hernández-Cano Alejandro, AlKhamissi Badr, Matoba Kyle, Salvi Francesco, Pagliardini Matteo, Fan Simin, Köpf Andreas, Mohtashami Amirkeivan, Sallinen Alexandre, Swamy Vinitra, Sakhaeirad Alireza, Krawczuk Igor, Bayazit Deniz, Marmet Axel, Li Mi, Boillat-Blanco Noémie, Keitel Kristina, Elkin Javier A., Robert Blaise, Montariol Syrielle, Bressan Silvia, Chen D., Demers Vincent, Emery N., Glasson Nicolas, Mensah Paulina Boadiwaa, Miauton Alix, Roemer S., Siebert Johan N, Starvaggi Carl, Suttels Véronique, Tan Rainer, Taylor Richard A., Toit Jacques du, Hartley Mary-Anne, Jaggi Martin, and Bosselut Antoine. Meditron: Open medical foundation models adapted for clinical practice. *Research Square (research Square)*, 2024. doi: 10.21203/rs.3.rs-4139743/v1.

[10] Qianqian Xie, Qingyu Chen, Aokun Chen, Cheng Peng, Yan Hu, Fongci Lin, Xueqing Peng, Jimin Huang, Jeffrey Zhang, Vipina Keloth, Xinyu Zhou, Huan He, Lucila Ohno-Machado, Yonghui Wu, Hua Xu, and Jiang Bian. Me-llama: Foundation large language models for medical applications. *Research Square*, 2024. ISSN 26935015. doi: 10.21203/rs.3.rs-4240043/v1.

[11] Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical images. *Nature Communications*, 15(1):654, 2024. ISSN 20411723. doi: 10.1038/s41467-024-44824-z.

[12] Joseph Cox, Peng Liu, Skylar E Stolte, Yunchao Yang, Kang Liu, Kyle B See, Huiwen Ju, and Ruogu Fang. Brainsegfounder: Towards 3d foundation models for neuroimage segmentation. *Arxiv*, 2024. ISSN 23318422.

[13] Zhihan Zhou, Yanrong Ji, Weijian Li, Pratik Dutta, Ramana Davuluri, and Han Liu. Dnabert-2: Efficient foundation model and benchmark for multi-species genome. 2024.

[14] Albi Celaj, Alice Jiexin Gao, Tammy T.Y. Lau, Erle M. Holgersen, Alston Lo, Varun Lodaya, Christopher B. Cole, and [Missing Name] Denroche. An rna foundation model enables discovery of disease mechanisms and candidate therapeutics. *bioRxiv*, 2023. doi: 10.1101/2023.09.20.558508. URL https://www.biorxiv.org/content/early/2023/09/26/2023.09.20.558508.

[15] Ruipu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Da Li, Pengcheng Lu, Tao Wang, Linmei Hu, Minghui Qiu, and Zhongyu Wei. Valley: Video assistant with large language model enhanced ability. 2023.

[16] Zhao Wang, Chang Liu, Shaoting Zhang, and Qi Dou. Foundation model for endoscopy video analysis via large-scale self-supervised pre-train. *Medical Image Computing and Computer Assisted Intervention – Miccai 2023*, pages 101–111, 2023. doi: 10.1007/978-3-031-43996-4_10.

[17] Sheng Zhang, Yanbo Xu, Naoto Usuyama, Hanwen Xu, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, Cliff Wong, Andrea Tupini, Yu Wang, Matt Mazzola, Swadheen Shukla, Lars Liden, Jianfeng Gao, Angela Crabtree, Brian Piening, Carlo Bifulco, Matthew P. Lungren, Tristan Naumann, Sheng Wang, and Hoifung Poon. Biomedclip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. 2025.

[18] Kun Yuan, Vinkle Srivastav, Tong Yu, Joël L Lavanchy, Jacques Marescaux, Pietro Mascagni, Nassir Navab, and Nicolas Padoy. Learning multi-modal representations by watching hundreds of surgical video lectures. *Medical Image Analysis*, 105:103644, 2025. ISSN 13618423, 13618415. doi: 10.1016/j.media.2025.103644.

[19] Michael Wornow, Ethan Steinberg, Rahul Thapa, Jason A. Fries, and Nigam H. Shah. Ehrshot: An ehr benchmark for few-shot evaluation of foundation models. *Advances in Neural Information Processing Systems*, 36, 2023. ISSN 10495258.

[20] Ethan Steinberg, Jason Fries, Yizhe Xu, and Nigam Shah. Motor: A time-to-event foundation model for structured medical records, 2023. URL https://arxiv.org/abs/2301.03150.

[21] Manuel Burger, Daphné Chopard, Malte Londschien, Fedor Sergeev, Hugo Yèche, Rita Kuznetsova, Martin Faltys, Eike Gerdes, Polina Leshetkina, Peter Bühlmann, and Gunnar Rätsch. A foundation model for intensive care: Unlocking generalization across tasks and domains at scale. *medRxiv*, 2025. doi: 10.1101/2025.07.25.25331635. URL https://www.medrxiv.org/content/early/2025/07/25/2025.07.25.25331635.

[22] Manuel Burger, Fedor Sergeev, Malte Londschien, Daphné Chopard, Hugo Yèche, Eike Gerdes, Polina Leshetkina, Alexander Morgenroth, Zeynep Babür, Jasmina Bogojeska, Martin Faltys, Rita Kuznetsova, and Gunnar Rätsch. Towards foundation models for critical care time series, 2024. URL https://arxiv.org/abs/2411.16346.

[23] Robin van de Water, Hendrik Schmidt, Paul Elbers, Patrick Thoral, Bert Arnrich, and Patrick Rockenschaub. Yet another icu benchmark: A flexible multi-center framework for clinical ml. *12th International Conference on Learning Representations, Iclr 2024*, 2024.

[24] Max Horn, Michael Moor, Christian Bock, Bastian Rieck, and Karsten Borgwardt. Set functions for time series. *37th International Conference on Machine Learning, Icml 2020*, 168147-6: 4303–4313, 2020.

[25] Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific Reports*, 8(1):6085, 2018. ISSN 20452322. doi: 10.1038/s41598-018-24271-9.

[26] Yicheng Luo, Bowen Zhang, Zhen Liu, and Qianli Ma. Hi-patch: Hierarchical patch GNN for irregular multivariate time series modeling, 2025. URL https://openreview.net/forum?id=OGtUfA6Amo.

[27] Hrayr Harutyunyan, Hrant Khachatrian, David C. Kale, Greg Ver Steeg, and Aram Galstyan. Multitask learning and benchmarking with clinical time series data. *Scientific Data*, 6(1):96, 2019. ISSN 20524463. doi: 10.1038/s41597-019-0103-9.

[28] Hugo Yèche, Rita Kuznetsova, Marc Zimmermann, Matthias Hüser, Xinrui Lyu, Martin Faltys, and Gunnar Rätsch. Hirid-icu-benchmark – a comprehensive machine learning benchmark on high-resolution icu data, 2022. URL https://arxiv.org/abs/2111.08536.

[29] Adibvafa Fallahpour, Mahshid Alinoori, Wenqian Ye, Xu Cao, Arash Afkanpour, and Amrit Krishnan. Ehrmamba: Towards generalizable and scalable foundation models for electronic health records. *Proceedings of Machine Learning Research*, 259:291–307, 2024. ISSN 26403498.

[30] Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multidimensional transformers. *arXiv preprint arXiv:1912.12180*, 2019.

[31] Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. Mimic-iii, a freely accessible critical care database. *Scientific Data*, 3(1):160035, 2016. ISSN 20524463. doi: 10.1038/sdata.2016.35.

[32] Alistair E. W. Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J. Pollard, Benjamin Moody, Brian Gow, Li-wei H. Lehman, Leo A. Celi, and Roger G. Mark. Mimic-iv, a freely accessible electronic health record dataset. *Scientific Data*, 10(1):1, 2023. ISSN 20524463. doi: 10.1038/s41597-022-01899-x.

[33] Tom J. Pollard, Alistair E. W. Johnson, Jesse D. Raffa, Leo A. Celi, Roger G. Mark, and Omar Badawi. The eicu collaborative research database, a freely available multi-center database for critical care research. *Scientific Data*, 5(1):180178, 2018. ISSN 20524463. doi: 10.1038/sdata.2018.178.

[34] Laurens Van Der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2625, 2008. ISSN 15337928, 15324435.

[35] Patrick Rockenschaub, Adam Hilbert, Tabea Kossen, Paul Elbers, Falk Von Dincklage, Vince Istvan Madai, and Dietmar Frey. The impact of multi-institution datasets on the generalizability of machine learning prediction models in the icu. *Critical Care Medicine*, 52(11):1710–1721, 2024. ISSN 15300293, 00903493. doi: 10.1097/CCM.0000000000006359.

[36] Weather data from the max planck institute for biogeochemistry. URL https://www.bgc-jena.mpg.de/wetter/.

[37] Artur Trindade. Electricity load diagrams, 2015. URL https://doi.org/10.24432/C58C86.

[38] R. van de Water. YAIB-cohorts, 2024. URL https://github.com/rvandewater/YAIB-cohorts. GitHub repository.

[39] Robin van de Water, Hendrik Schmidt, and Patrick Rockenschaub. Yet Another ICU Benchmark (YAIB). URL https://github.com/rvandewater/YAIB.

[40] Nicolas Bennett, Drago Plečko, Ida Fong Ukor, Nicolai Meinshausen, and Peter Bühlmann. ricu: R's interface to intensive care data. *Gigascience*, 12:giad041, 2023. ISSN 2047217x. doi: 10.1093/gigascience/giad041.

[41] Sindhu Tipirneni and Chandan K. Reddy. Self-supervised transformer for sparse and irregularly sampled multivariate clinical time-series. *Acm Transactions on Knowledge Discovery From Data*, 16(6):105, 2022. ISSN 1556472x, 15564681. doi: 10.1145/3516367.

[42] Hrishikesh Patel, Ruihong Qiu, Adam Irwin, Shazia Sadiq, and Sen Wang. Emit- event-based masked auto encoding for irregular time series. *Arxiv (cornell University)*, 2024. doi: 10.48550/arXiv.2409.16554.

[43] Alex Labach, Aslesha Pokhrel, Xiao Shi Huang, Saba Zuberi, Seung Eun Yi, Maksims Volkovs, Tomi Poutanen, and Rahul G. Krishnan. Duett: Dual event time transformer for electronic health records. 2023.

[44] Alistair Johnson, Tom Pollard, and Roger Mark. MIMIC-III Clinical Database CareVue subset (version 1.4). https://doi.org/10.13026/8a4q-w170, 2022. PhysioNet. RRID:SCR_007345.

A Appendix

A.1 BAT architecture

One Transformer-based model designed for irregular multivariate time series data is Bi-Axial Transformer (BAT) introduced in the work of DeVries et al. [2]. Unlike classical Transformers that use an encoder-decoder structure for sequence generation, BAT relies solely on encoder-style self-attention to model dependencies across both the temporal and feature/sensor dimensions. This axial attention structure allows BAT to capture rich interactions within and across time steps and modalities. To address common challenges in irregular multivariate time series data, such as missing values and heterogeneous input types, BAT incorporates missingness indicators directly into its input embeddings. Figure 2 (Adapted from Figure 2 in DeVries et al. [2]) illustrates the overall architecture of BAT. The input to the model is a multivariate time series of shape $D \times T$, where D is the number of features and T is the number of time steps. Each input is embedded using a combination of the observed value (including missingness), a learned feature identity embedding, and a continuous time-based positional encoding. This results in an embedded tensor of shape $D \times T \times E$, where E is the embedding dimension. The embedded input is then processed by bi-axial attention layers, which apply self-attention separately along the time and feature dimensions. The output from the attention layers is subsequently pooled and concatenated with the static features, P, which include non-time-varying demographic variables such as age and sex. The BAT model was adapted in this work to have two different prediction heads: one for a supervised learning setup via binary classification, predicting $\hat{y} \in \{0, 1\}$, and another for a self-supervised learning setup via forecasting, predicting $\hat{\mathbf{X}}^{\text{for}} \in \mathbb{R}^{T^{\text{for}} \times D}$.

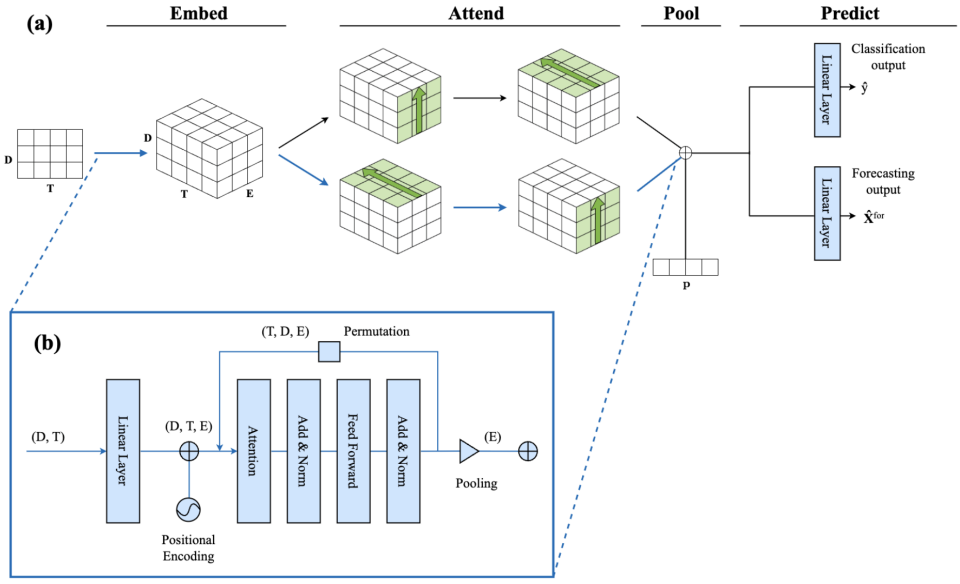

Figure 2: Overview of the Bi-Axial Transformer (BAT) architecture. (a) Shows the full model architecture and data representation, and (b) Shows an attention track indicated by the blue arrows. D is the number of time-varying features, T the number of time steps, E the embedding size, and P the static features. The model supports both binary classification and forecasting via separate prediction heads. This figure is adapted from Figure 2 in DeVries et al. [2].

A.2 Yet another ICU benchmark framework

The Yet Another ICU Benchmark (YAIB) framework [23] provides a modular, end-to-end solution for clinical machine learning on ICU data, explicitly designed to address key limitations regarding reproducibility in the field. An illustration of the framework is shown in Figure 3 (Inspired by Figure 1 in van de Water et al. [23]). It consists of two repositories: `YAIB-cohorts` [38], which handles dataset harmonization and cohort construction, and the main `YAIB` repository [39], which manages model training and evaluation. The `YAIB-cohorts` repository builds on the open-source `R` package `ricu` [40] to harmonize multiple ICU datasets using a unified, concept-based abstraction of clinical variables. It supports five publicly available EHR datasets (MIMIC-III, MIMIC-IV, eICU, HiRID, AUMCdb) and maps their contents into a common structure with consistent semantic definitions and temporal alignment, and supports integration of new datasets. This harmonization enables standardized cohort construction, label definition, and data extraction across datasets, facilitating multi-center analyses and reproducible experimental setups. The main `YAIB` repository then provides the downstream machine learning pipeline for supervised modeling, including pre-processing, feature extraction, and evaluation. We entended the framework to support self-supervised pretraining, described in detail in Appendix A.3, as well as add the BAT architecture to the selection of models. The models in this work show lower performance on the mortality prediction task compared to similar studies [41, 42, 43], particularly in class sensitive metric, AUC-PR. As van de Water et al. [23] notes, variations in preprocessing pipelines, task and cohort definitions etc. and limited transparency hinder reproducibility, especially when code is unavailable. In contrast to most prior work that experiment with performance-boosting strategies, e.g. upsampling of the minority class, our main focus of this work has been on demonstrating the transfer learning potential of robust self-supervised models rather than surpassing state-of-the-art results. We used a weighted loss available in the YAIB framework to handle the high class imbalance, but did not investigate further performance-boosting strategies. This work benchmarks pre-trained models against two baselines, aiming to provide a comprehensive and fair comparison throughout the entire machine learning pipeline.

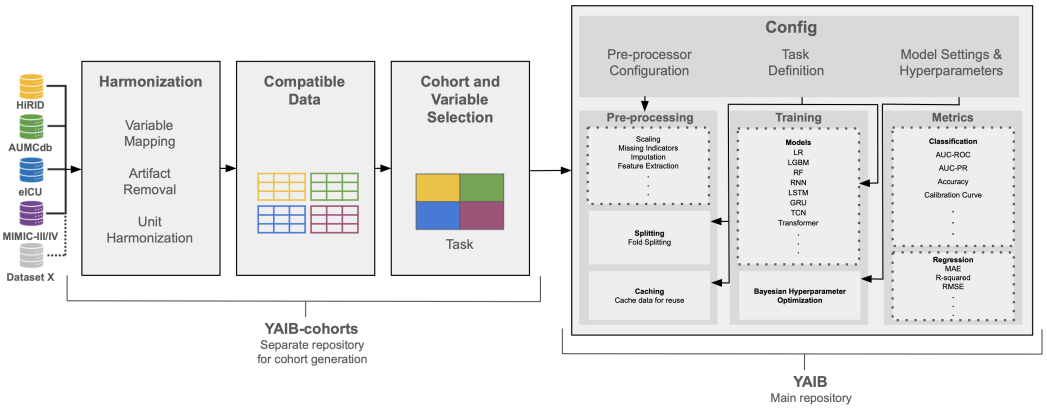

Figure 3: Overview of the Yet Another ICU Benchmark (YAIB) pipeline. The left side illustrates the creation of harmonized ICU cohorts, implemented in a separate repository, `YAIB-cohorts` [38]. The right side represents the machine learning component of the pipeline, contained in the main `YAIB` repository [39], which covers preprocessing, model training, and evaluation. Dotted-line components indicate extensible modules that follow a standardized interface. This figure is inspired from Figure 1 in van de Water et al. [23].

A.3 Implementation of self-supervised learning

The self-supervised learning objective is implemented as a forecasting task, where future values are predicted based on past values. The approach in this work is inspired by the dynamic sampling method proposed by Tipirneni and Reddy [41], in which observation and forecasting windows are dynamically constructed during batch loading. A set of constraints is introduced to govern the selection of observation and forecasting windows. These constraints are designed to ensure both sufficient historical context and generating a valid forecasting window, thereby improving the overall quality and consistency of the training data. The overall goal is to select a valid time index where the observation and forecasting window are split. This work performs the search for the valid time index in a batch-wise manner. For each batch element, a patient is randomly sampled and candidate time indices are identified. These indices are filtered based on the following constraints:

1. **Sparsity check:** The selected index must correspond to a time point where there is at least one observed value in the observation window.
2. **Minimum Observation Length:** The index must be at least L time steps into the time-series to ensure sufficient historical context. For this work, $L = 12$ hours.
3. **Forecasting Window Availability:** The index must allow room for a complete forecasting window of length H. For this work, $H = 2$ hours.

If valid time indices are found, one is randomly selected as the split point between the observation window and the forecasting window. If no valid index is found, a new patient is sampled and the process is repeated. This approach ensures that the model is trained on different points within each hospital stay, with varying observation window lengths, effectively exposing it to diverse temporal contexts and clinical information.The sampling method are defined in Algorithm 1.

Algorithm 1 Dynamic sampling of observation and forecasting window during batch loading

function CALL(batch)
  $(data, mask) \leftarrow$ LOADBATCH$(batch)$
  $(B, C, T) \leftarrow$ SHAPE$(data)$
  $t_1 \leftarrow$ None, $tries \leftarrow 0$, $max\_tries \leftarrow B$
  while $t_1 =$ None **and** $tries < max\_tries$ **do**
    $i \leftarrow$ RANDOMINTEGER$(0, B-1)$
    $valid\_times \leftarrow \{t \mid mask[i,t] = \text{True}\}$
    $valid\_times \leftarrow \{t \in valid\_times \mid t \geq L\}$
    if $valid\_times = \emptyset$ **then**
      $tries \leftarrow tries + 1$
      continue
    end if
    $max\_index \leftarrow \max(valid\_times)$
    $valid\_times \leftarrow \{t \in valid\_times \mid t \leq max\_index - H\}$
    if $valid\_times = \emptyset$ **then**
      $tries \leftarrow tries + 1$
      continue
    end if
    $t_1 \leftarrow$ RANDOMCHOICE$(valid_times)$
  end while
  if $t_1 =$ None **then**
    raise Error("No valid index found in batch")
  end if
  $t_0 \leftarrow \max(0,\ t_1 - max\_obs)$
  $t_2 \leftarrow t_1 + forecast\_horizon$
  $X_{obs} \leftarrow data[:, :, t_0 : t_1]$
  $M_{obs} \leftarrow mask[:, :, t_0 : t_1]$
  $X_{for} \leftarrow data[:, :, t_1 : t_2]$
  $M_{for} \leftarrow mask[:, :, t_1 : t_2]$
  return $(X_{obs}, M_{obs}, X_{for}, M_{for})$
end function

A.4 Datasets, cohort definitions and preprocessing

The three publicly available ICU datasets used in this work are MIMIC-III [31], MIMIC-IV [32], and the eICU [33]. These datasets contain structured clinical data from ICU stays and are widely used in the development of machine learning models for critical care. Table 2 summarizes key statistics for the three datasets. As this study involves pooling of the datasets into one larger dataset, it is important to avoid patient overlap between them, as this could lead to data leakage in downstream experiments. Therefore, a filtered subset of the original MIMIC-III dataset that excludes all patients also found in MIMIC-IV is used. The overlap exists because MIMIC-IV was developed to extend MIMIC-III by including more recent and higher-resolution data. The subset of the original MIMIC-III dataset used in this work is called MIMIC-III Clinical Database CareVue subset [44], which limits MIMIC-III to records from 2001–2008 thereby excluding patient stays in the overlapping time period. Furthermore, a second filter was applied to remove any individual who also appears in MIMIC-IV, as some patients had an earlier stay recorded in MIMIC-III and a later hospital stay captured in MIMIC-IV. Feature

Table 2: Dataset statistics for MIMIC-III, MIMIC-IV, and eICU. BIDMC = Beth Israel Deaconess Medical Center. *The subset of the MIMIC-III data used in this work is the MIMIC-III CareVue subset [44]; values in parentheses represent statistics from the full MIMIC-III dataset.

Dataset	**Admissions**	**Collection period**	**Origin**	**Hospital**	**Mortality (Positive class)**
MIMIC-III*	27k* (40k)	2001–2008* (2001–2012)	U.S.	BIDMC	11.9%
MIMIC-IV	73k	2008–2019	U.S.	BIDMC	7.3%
eICU	200k	2014–2015	U.S.	208 hospitals across the U.S.	5.5%

selection was performed by van de Water et al. [23] based on availability across all benchmarked datasets available in YAIB [39], with an emphasis on consistency to support cross-dataset experiments. A total of 52 clinical features were used as input for model development, consisting of 4 static and 48 time-varying variables. A full list of the features used and their units are provided in Table 3.

Patient cohorts were constructed using the `YAIB-cohorts` repository [38]. Data were temporally aligned and resampled at one-hour resolution. Uniform exclusion criteria were applied across all datasets and tasks: **(1)** invalid ICU stay timing (e.g., negative length of stay), **(2)** ICU stay < 6 hours, **(3)** fewer than four valid time points, **(4)** measurement gaps > 12 hours, and **(5)** age < 18 years at ICU admission. These filters preceded task-specific cohort definitions. For mortality classification, the input window was the first 24 hours post-ICU admission; stays < 30 hours were excluded to prevent causal leakage [23]. The outcome label was mortality = 1 if the patient died during the same hospital admission. The YAIB framework requires task-specific definitions with inclusion/exclusion criteria and label generation, making it incompatible with self-supervised learning where labels are not used. Since YAIB cannot generate unlabeled cohorts, we selected the Length of Stay task, which imposed no additional exclusions beyond the five main criteria mentioned, thereby maximizing size of the pre-training dataset. Length of stay labels were still generated as a consequence of YAIB's supervised design but were simply discarded resulting in a unlabeled cohort generation. This unlabeled cohort was the patient cohort used for pre-training. This approach maintained full compatibility with the existing framework while avoiding the need for significant changes to its design.

Preprocessing was carried out using YAIB's main repository `YAIB` [39], applied after cohort extraction. This included the addition of missingness indicators, forward-fill imputation within each ICU stay, and mean imputation for values without prior observations, using statistics computed from the training set to avoid data leakage. All features were standardized to zero mean and unit variance based on training split statistics. Data were split into training, validation, and test sets using cross-validation folds defined during preprocessing, ensuring that preprocessing steps such as scaling and imputation were fitted only on the training data. All preprocessing decisions followed the default setup specified by the `base_classification_preprocessor` class in the main repository, `YAIB`.

Table 3: Clinical features and units

Feature	Unit
Static	
Age at hospital admission	years
Female sex	–
Patient height	cm
Patient weight	kg
Time-varying	
Albumin	g/dL
Alkaline phosphatase	IU/L
Alanine aminotransferase	IU/L
Aspartate aminotransferase	IU/L
Band form neutrophils	%
Base excess	mmol/L
Bicarbonate	mmol/L
Bilirubin (direct)	mg/dL
Bilirubin (total)	mg/dL
Blood pressure (diastolic)	mmHg
Blood pressure (systolic)	mmHg
Blood urea nitrogen	mg/dL
Calcium	mg/dL
Calcium ionized	mmol/L
Chloride	mmol/L
CO_2 partial pressure	mmHg
C-reactive protein	mg/L
Creatinine	mg/dL
Creatine kinase	IU/L
Creatine kinase MB	ng/mL
Fibrinogen	mg/dL
Fraction of inspired oxygen	%
Glucose	mg/dL
Haemoglobin	g/dL
Heart rate	beats/minute
International normalised ratio (INR)	–
Lactate	mmol/L
Lymphocytes	%
Magnesium	mg/dL
Mean arterial pressure	mmHg
Mean cell haemoglobin	pg
Mean corpuscular haemoglobin concentration	%
Mean corpuscular volume	fL
Methaemoglobin	%
Neutrophils	%
O_2 partial pressure	mmHg
Oxygen saturation	%
Partial thromboplastin time	sec
pH of blood	–
Phosphate	mg/dL
Platelets	$1{,}000/\mu$L
Potassium	mmol/L
Respiratory rate	breaths/minute
Sodium	mmol/L
Temperature	°C
Troponin T	ng/mL
Urine output	mL
White blood cells	$1{,}000/\mu$L

A.5 Data distributions and transferability

A t-SNE [34] analysis of the preprocessed and harmonized datasets was performed to asses the data distributions of the three datasets before modeling. The results shown in Figure 4. It reveals that the datasets appear to lie within the same general distribution, which might be a result of similar clinical practices and patient populations. BAT was used to train three models, each trained and hyperparameter tuned on one of the three datasets. The models were then tested independently on all datasets. The models were trained in the 5-fold cross-validation setup described in Appendix A.7. The average and standard deviations of AUC-ROC and AUC-PR can be seen in Figure 5. Each model achieves the best performance on the dataset it was trained on (diagonal). Performance drops relative to the diagonal performance is highest for the model trained on MIMIC-III (AUC-ROC drops between 4.1–4.8) and the lowest for the model trained on MIMIC-IV (AUC-ROC drops between 1.0-1.8) showing different degrees of transferability among the datasets.

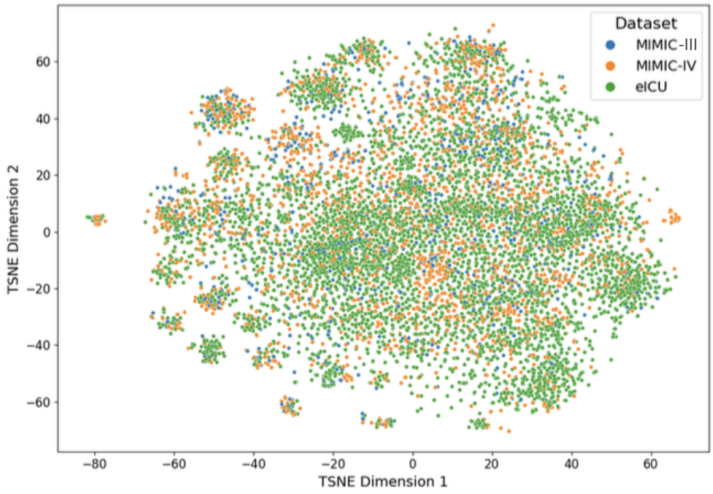

Figure 4: Two-dimensional t-SNE projection [34] of the harmonized and preprocessed datasets used in this study: MIMIC-III, MIMIC-IV, and eICU. Each point represents a time step of a patient stay.

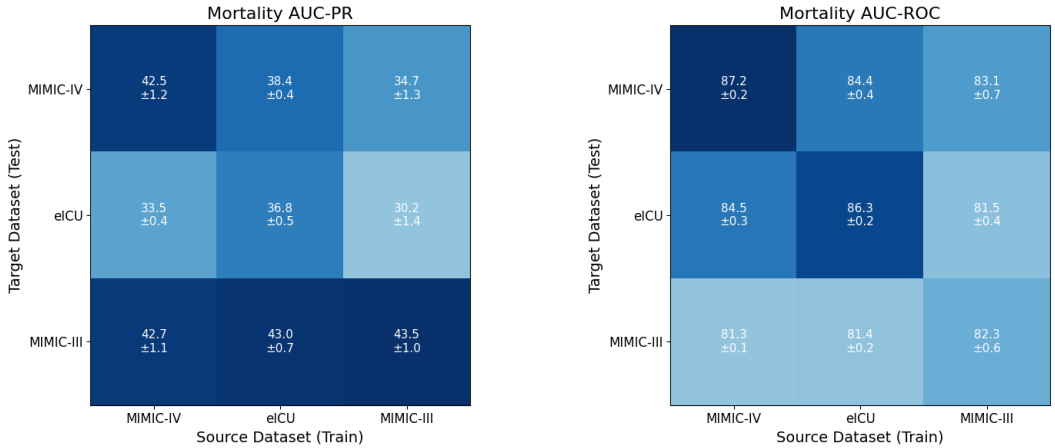

Figure 5: Performance of three independently trained models, each evaluated on all three datasets: MIMIC-III, MIMIC-IV, and eICU. **(a)** AUC-PR scores; **(b)** AUC-ROC scores. Values reflect mean performance ± standard deviation on held-out test sets using 5-fold cross-validation.

A.6 Hyperparameters

This appendix provides the optimal hyperparameters used across all experiments. Hyperparameters were selected for each model in two stages. First, a grid search was used to identify a suitable hyperparameter range. Once the hyperparameter ranges were identified, Bayesian hyperparameter optimization (implemented by van de Water et al. [39] in the YAIB framework) was used to select the final hyperparameters. During fine-tuning the models batch size and learning rate were fine-tuned using a grid search while the rest of the hyperparameters were kept the same as during pre-training.

Table 4: Final hyperparameters for pretraining of BAT on pooled datasets MIMIC-III + MIMIC-IV.

Component	Hyperparameters
Model	attn_dropout = 0.207, dropout = 0.364, heads = 1, layers = 2, pooling = max, use_mask = False, value_embed_size = 128
Trainer	batch_size = 64, epochs = 200, patience = 15, min_delta = 5e-3
Optimizer	lr = 7.781e-4, weight_decay = 1e-6
Forecasting	forecast_horizon = 2, sensors_count = 48

Table 5: Final hyperparameters for pretraining of BAT on pooled datasets eICU + MIMIC-III.

Component	Hyperparameters
Model	attn_dropout = 0.357, dropout = 0.249, heads = 1, layers = 8, pooling = max, use_mask = False, value_embed_size = 64
Trainer	batch_size = 64, epochs = 200, patience = 10, min_delta = 5e-3
Optimizer	lr = 6.196e-4, weight_decay = 1e-6
Forecasting	forecast_horizon = 2, sensors_count = 48

Table 6: Final hyperparameters for pretraining of BAT on pooled datasets eICU + MIMIC-IV.

Component	Hyperparameters
Model	attn_dropout = 0.284, dropout = 0.348, heads = 1, layers = 6, pooling = max, use_mask = False, value_embed_size = 64
Trainer	batch_size = 64, epochs = 200, patience = 10, min_delta = 5e-3
Optimizer	lr = 3.375e-4, weight_decay = 1e-6
Forecasting	forecast_horizon = 2, sensors_count = 48

Table 7: Final hyperparameters for fine-tuning the classification head of BAT on the datasets MIMIC-III, MIMIC-IV and eICU. All other hyperparameters are kept the same as during pre-training.

Dataset	Hyperparameters
MIMIC-III	batch_size = 64, learning_rate = 1e-2, lr_scheduler = ExponentialLR(gamma = 0.95)
MIMIC-IV	batch_size = 64, learning_rate = 5e-3 , lr_scheduler = ExponentialLR(gamma = 0.95)
eICU	batch_size = 24, learning_rate = 7e-3, lr_scheduler = ExponentialLR(gamma = 0.95)

Table 8: Final hyperparameters, for fine-tuning the full model, BAT on the datasets MIMIC-III, MIMIC-IV and eICU. All other hyperparameters are kept the same as during pre-training.

Dataset	**Hyperparameters**
MIMIC-III	batch_size = 64, learning_rate = 9e-05, lr_scheduler = ExponentialLR(gamma = 0.95)
MIMIC-IV	batch_size = 24, learning_rate = 7e-4, lr_scheduler = ExponentialLR(gamma = 0.95)
eICU	batch_size = 64, learning_rate = 5e-4 , lr_scheduler = ExponentialLR(gamma = 0.95)

A.7 Data splits

Models in Figure 5 were trained using a 5-fold cross-validation setup. An initial 80/20 split separated the data into training and test sets, and the training set was further divided into five folds, rotating the validation fold across training iterations. The final performance of the models are reported as the mean ± standard deviation across the five cross-validated models. Pre-training of the multi-dataset model was done using the same cross-validation setup. However, only one of the five models was used for the zero-shot and fine-tuning experiments. This was the one that achieved the lowest masked mean squared error loss during pre-training. The baseline models in the fine-tuning experiment used the same data split as the selected pre-trained model.

A.8 Models and code availability

All code related to this project is available at: `https://anonymous.4open.science/r/YAIB-04BC/README.md` . This repository was originally forked from `https://github.com/rvandewater/YAIB` [39] and has been expanded to support the experiments and methods presented in this work.

Table 9: Pre-trained models on combinations of the three datasets of this work (MIMIC-III, MIMIC-IV, and eICU). Details on the models, including training and computational resources, are summarized.

	Pre-trained Models		
Pooled pre-training datasets	MIMIC-IV + eICU	MIMIC-III + eICU	MIMIC-III + MIMIC-IV
Fine-tuning dataset	MIMIC-III	MIMIC-IV	eICU
#Params	0.86M	1.13M	0.97M
Pre-training dataset size	273k	227k	100k
Pre-training positive class	6.0%	6.26%	8.54%
Fine-tuning positive class	11.9%	7.3%	5.5%
Pre-training time	$\sim$ 7h	$\sim$ 7h	$\sim$ 3h
Pre-training GPU	1 x A100 (40 GB)	1 x A100 (40 GB)	1 x A100 (40 GB)
Fine-tuning time	$\leq$ 20 min	$\leq$ 20 min	$\leq$ 20 min
Fine-tuning GPU	1 x V100 (16 GB)	1 x V100 (16 GB)	1 x V100 (16 GB)

A.9 Fine-tuning results

Table 11: Model performance, mean AUC-PR ± sd , across dataset size ranging from 100 to 9,506 on MIMIC-III, MIMIC-IV, and eICU. The pre-trained BAT models are fine-tuned and the baseline models are trained from scratch on the subsets. Highest performance for each dataset size is indicated in bold, and second highest is underlined.

Fine-tuning/ Training dataset	Dataset size	BAT (Fine-tuned full)	BAT (Fine-tuned head)	BAT (Scratch)	Transformer (Scratch)
MIMIC-III	100	9.73 ± 0.02	**25.77 ± 5.85**	15.40 ± 2.73	11.94 ± 0.76
	500	8.68 ± 0.03	**30.36 ± 1.52**	23.98 ± 2.62	16.75 ± 4.65
	1000	**36.24 ± 1.63**	33.98 ± 1.32	27.63 ± 2.23	21.34 ± 4.58
	2000	**39.01 ± 1.12**	37.50 ± 0.78	30.26 ± 3.07	25.41 ± 0.65
	3000	**40.27 ± 1.14**	38.35 ± 1.62	31.92 ± 1.48	26.10 ± 1.18
	5000	**41.89 ± 1.31**	38.99 ± 0.96	36.30 ± 0.31	26.14 ± 0.40
	7000	**42.55 ± 1.54**	40.22 ± 0.54	36.32 ± 1.43	26.40 ± 0.67
	9000	**43.57 ± 0.94**	40.14 ± 0.70	37.09 ± 1.06	27.13 ± 0.41
	9560	**43.97 ± 0.00**	40.78 ± 0.00	36.73 ± 0.00	27.86 ± 0.00
MIMIC-IV	100	**12.50 ± 4.43**	4.90 ± 0.02	8.92 ± 4.98	7.72 ± 0.66
	500	**23.96 ± 3.07**	23.88 ± 0.96	22.63 ± 3.23	10.92 ± 2.65
	1000	**28.98 ± 0.85**	26.97 ± 1.71	26.12 ± 1.95	13.06 ± 1.56
	2000	**33.19 ± 2.95**	28.99 ± 1.38	28.79 ± 1.64	16.47 ± 0.70
	3000	**36.60 ± 1.69**	30.23 ± 1.24	32.21 ± 2.63	16.86 ± 0.59
	5000	**38.10 ± 1.36**	31.31 ± 1.17	34.97 ± 1.03	18.00 ± 1.11
	7000	**38.60 ± 1.77**	31.92 ± 0.60	37.68 ± 1.25	18.65 ± 0.66
	9000	**38.75 ± 0.53**	32.19 ± 1.00	38.41 ± 0.88	18.91 ± 1.36
	9560	**40.05 ± 0.80**	32.16 ± 1.00	38.75 ± 1.84	19.62 ± 0.97
eICU	100	7.24 ± 0.09	**14.20 ± 2.77**	10.79 ± 1.83	8.82 ± 1.26
	500	**23.54 ± 3.20**	23.38 ± 0.98	16.78 ± 2.07	8.13 ± 0.87
	1000	**28.37 ± 1.13**	25.39 ± 1.56	20.86 ± 3.31	6.58 ± 4.00
	2000	**31.89 ± 0.77**	27.46 ± 0.50	26.12 ± 0.86	13.71 ± 1.03
	3000	**32.32 ± 1.07**	28.43 ± 0.70	28.00 ± 1.38	13.90 ± 0.99
	5000	**33.89 ± 0.49**	29.09 ± 0.62	30.13 ± 1.37	14.41 ± 0.42
	7000	**35.14 ± 0.33**	29.39 ± 0.77	31.09 ± 0.98	13.98 ± 0.49
	9000	**35.20 ± 0.87**	29.53 ± 0.95	31.59 ± 0.94	14.61 ± 1.06
	9506	**35.17 ± 0.54**	29.40 ± 1.18	31.67 ± 1.25	14.18 ± 1.25

Table 10: Model performance, mean AUC-ROC ± sd , across dataset size ranging from 100 to 9,506 on MIMIC-III, MIMIC-IV, and eICU. The pre-trained BAT models are fine-tuned and the baseline models are trained from scratch on the subsets. Highest performance for each dataset size is indicated in bold, and second highest is underlined.

Fine-tuning/ Training dataset	Dataset size	BAT (Fine-tuned full)	BAT (Fine-tuned head)	BAT (Scratch)	Transformer (Scratch)
MIMIC-III	100	40.88 ± 0.12	**66.14 ± 5.32**	57.37 ± 6.33	48.91 ± 2.23
	500	35.27 ± 0.18	72.14 ± 1.64	**72.60 ± 1.91**	56.52 ± 7.43
	1000	**78.30 ± 0.66**	74.98 ± 1.23	74.35 ± 0.89	62.83 ± 7.95
	2000	**79.72 ± 0.57**	78.05 ± 0.63	75.84 ± 1.98	69.99 ± 0.56
	3000	**80.43 ± 0.61**	78.42 ± 1.15	77.73 ± 0.69	71.17 ± 0.44
	5000	**81.22 ± 0.63**	78.84 ± 0.55	79.77 ± 0.69	71.20 ± 0.46
	7000	**81.84 ± 0.71**	79.85 ± 0.31	80.10 ± 0.20	72.47 ± 0.30
	9000	**82.50 ± 0.41**	79.97 ± 0.35	81.12 ± 0.60	72.42 ± 0.39
	9560	**82.37 ± 0.00**	80.10 ± 0.00	80.73 ± 0.00	72.99 ± 0.00
MIMIC-IV	100	**58.04 ± 7.67**	30.67 ± 0.19	49.62 ± 11.54	47.82 ± 2.80
	500	74.85 ± 3.56	70.51 ± 1.18	**75.17 ± 2.82**	56.21 ± 7.67
	1000	**78.85 ± 1.11**	73.52 ± 2.04	77.84 ± 0.34	60.74 ± 3.52
	2000	**81.63 ± 2.01**	76.57 ± 0.88	79.81 ± 1.15	69.33 ± 1.24
	3000	**83.61 ± 1.24**	77.87 ± 1.61	81.97 ± 1.79	69.43 ± 1.74
	5000	**84.67 ± 0.81**	79.43 ± 0.40	83.76 ± 0.83	72.93 ± 0.66
	7000	**84.91 ± 0.93**	80.16 ± 0.30	84.69 ± 0.66	73.63 ± 0.64
	9000	**85.39 ± 0.59**	80.60 ± 0.51	84.83 ± 0.49	73.51 ± 1.30
	9560	**85.28 ± 0.80**	80.70 ± 0.43	85.03 ± 0.52	73.73 ± 0.77
eICU	100	47.87 ± 0.23	64.11 ± 3.81	**65.51 ± 3.07**	60.43 ± 5.30
	500	**77.80 ± 2.00**	71.65 ± 1.36	72.79 ± 1.68	56.09 ± 4.40
	1000	**81.42 ± 0.74**	74.22 ± 1.14	75.74 ± 2.08	47.93 ± 11.66
	2000	**82.99 ± 1.00**	76.24 ± 0.99	79.46 ± 0.98	69.51 ± 0.80
	3000	**83.03 ± 0.50**	76.68 ± 0.99	81.17 ± 0.47	69.69 ± 0.30
	5000	**84.28 ± 0.33**	77.89 ± 1.17	82.25 ± 0.56	70.26 ± 1.02
	7000	**84.66 ± 0.32**	78.38 ± 0.63	83.25 ± 0.35	70.49 ± 0.81
	9000	**84.85 ± 0.58**	78.80 ± 1.04	83.43 ± 0.21	71.08 ± 0.78
	9506	**85.05 ± 0.25**	78.26 ± 1.00	83.48 ± 0.58	71.10 ± 0.61