ORIGINAL RESEARCH Open access

Meta-Learning Framework for Rapid Adaptation of Sepsis Prediction Models across Different Intensive Care Units with Varying Data Availability and Patient Demographics

Ali Hassan[1*], Noor Siddiqui[1], Bilal Khan[2], Sana Malik[1]

Abstract

Sepsis prediction models perform poorly when transferred between ICUs due to demographic and practice variation, leading to substantial performance drops caused by differences in patient populations, admission criteria, and data recording standards, which limits reliable deployment across healthcare systems. Retraining models from scratch requires large labeled datasets that many ICUs lack due to cost, time, and resource limitations, making it difficult for low-resource settings to develop or adopt effective predictive tools. We propose a meta-learning approach that enables rapid adaptation of sepsis prediction models using few-shot local data, leveraging pre-training across multiple ICUs to support fast personalization in new environments. The framework includes meta-training across diverse source ICUs to learn a generalizable initialization and meta-adaptation at the target ICU using only a few gradient updates on limited data, enabling efficient few-shot learning. This approach improves sepsis prediction in low-resource and heterogeneous ICU settings by reducing data requirements and increasing robustness to demographic shifts, supporting more equitable access to AI tools in critical care. The proposed framework enables efficient and fair deployment of sepsis prediction models across diverse ICUs, bridging resource gaps and improving scalability and adaptability of clinical AI systems globally.

Keywords Domain adaptation, Sepsis prediction, Few-shot learning, Meta-learning, Cross-ICU adaptation, Demographic shift

*Correspondence:
Ali Hassan
ali.hassan@gmail.com
[1] Department of AI in Clinical Systems, Aga Khan University, Karachi, Pakistan
[2] Department of Intelligent Healthcare Engineering, Qatar University, Doha, Qatar

Introduction

Sepsis prediction models developed on large academic ICUs fail at community hospitals; performance drops due to case mix, demographics, data quality [1, 2]. Sepsis prediction models developed on large academic ICUs frequently fail when applied to community hospitals due to substantial differences in case mix, patient demographics, and data quality [1, 2]. These discrepancies lead to degraded predictive performance that can compromise clinical decision-making [3, 4]. Domain adaptation techniques have been investigated to mitigate these issues but often require additional resources [5]. Furthermore, variations in electronic health record systems exacerbate the generalization problems [6]. Therefore, new frameworks are needed to ensure robust performance across all ICU types [7].

Collecting large labeled datasets at every ICU is infeasible (cost, expertise, time); low-resource ICUs left behind [3, 4].

Collecting large labeled datasets at every ICU is infeasible due to prohibitive costs, required expertise, and extended time frames [8, 9]. Low-resource ICUs are consequently left behind in the adoption of state-of-the-art sepsis prediction tools [10, 11]. This inequity highlights the unsustainability of traditional model development approaches [12]. Studies emphasize that data collection burdens hinder widespread AI implementation in critical care [13]. Innovative methods that bypass extensive labeling are thus critical for inclusive healthcare AI [6].

Meta-learning enables rapid adaptation with few examples; model learns how to learn across diverse ICUs [1, 5]. Meta-learning enables rapid adaptation with few examples by allowing models to learn how to learn effectively across diverse ICUs [14, 15]. This capability is particularly relevant for healthcare applications characterized by data heterogeneity [16]. The paradigm shifts the focus from task-specific training to meta-knowledge acquisition [17]. It has demonstrated potential in clinical risk prediction scenarios with limited samples [18]. As such, meta-learning offers a transformative solution for cross-ICU challenges [19].

This paper proposes a meta-learning framework for sepsis prediction that adapts to new ICU with ≤100 labeled patients; roadmap. This article presents a meta-learning framework for sepsis prediction that adapts to a new ICU with no more than 100 labeled patients and outlines a complete implementation roadmap. The proposed system leverages meta-training to create adaptable initializations suitable for few-shot scenarios [20, 21]. It provides detailed guidance on both training and adaptation phases to facilitate practical deployment [22]. The framework integrates seamlessly with existing clinical workflows while addressing key limitations of prior methods [23]. This thesis establishes a foundation for equitable and efficient AI use in intensive care units worldwide [24].

Background

Domain shift in sepsis prediction

Domain shift in sepsis prediction manifests as significant performance degradation when models are transferred between different ICUs. The primary sources include variations in patient demographics, admission criteria, practice patterns, and data coding conventions [1, 2, 3]. These factors create distributional mismatches that undermine the reliability of predictions in new settings [4, 10]. Research has consistently shown that unaddressed domain shift leads to increased false negatives in sepsis detection [11]. Mitigating these effects requires advanced adaptation strategies tailored to critical care data [5].

Additional contributors to domain shift encompass differences in clinician documentation styles and hospital resource levels [6, 13]. In sepsis contexts, such shifts can alter the prevalence and presentation of the condition across ICUs [7, 8]. Quantitative analyses reveal that these heterogeneities are more pronounced than in other medical domains [9]. Consequently, standard models require specialized handling to maintain efficacy [12]. This section establishes the foundational problem that the meta-learning framework seeks to resolve.

Transfer learning and domain adaptation

Transfer learning has been widely explored as a strategy to adapt sepsis prediction models between ICUs through techniques such as fine-tuning. Domain adversarial networks have also been employed to align feature distributions across different hospital datasets [2, 10]. These methods aim to reduce the discrepancy between source and target domains in clinical prediction tasks [4, 6]. However, they typically require a moderate amount of labeled data from the target ICU to achieve satisfactory performance [11]. Despite their utility, conventional transfer learning approaches often struggle with extreme domain shifts common in critical care [12].

Limitations of these techniques become evident when data availability is severely restricted or when demographic differences are pronounced [7, 8]. Fine-tuning alone may lead to overfitting on small target datasets without capturing the broader meta-knowledge [9]. Domain adaptation methods, while effective in some scenarios, do not fully address the rapid personalization needs of new ICUs [3]. Recent advancements have attempted to combine adversarial training with other regularization strategies [13]. Nonetheless, these still fall short for few-shot scenarios prevalent in low-resource settings [5].

Few-shot learning

Few-shot learning techniques, including prototypical networks and matching networks, offer potential solutions for scenarios with limited labeled data in healthcare. These methods enable models to classify or predict based on a small number of examples by learning effective similarity metrics [14, 15]. In the context of clinical applications, they have been applied to various diagnostic tasks with promising initial results [16]. However, their extension to time series data typical of ICU monitoring remains challenging due to the sequential nature of patient records [17]. Adapting these to sepsis prediction requires careful consideration of temporal dependencies [19].

Despite their data efficiency, few-shot learning approaches exhibit limitations when dealing with the high-dimensional and noisy characteristics of clinical time series [20, 21]. Prototypical networks may not sufficiently capture the complex dynamics of sepsis progression across heterogeneous patient populations [22]. Matching networks similarly struggle with the variability in ICU data formats [23]. These constraints highlight the need for more specialized few-shot strategies tailored to critical care [24]. Integrating few-shot principles with meta-learning could overcome some of these shortcomings [25].

Model-Agnostic Meta-Learning (MAML)

Model-Agnostic Meta-Learning (MAML) provides a powerful framework for rapid task adaptation through its distinctive inner and outer loop optimization structure. In the inner loop, task-specific parameters are updated via a few gradient steps on the support set for each ICU task [16]. This allows the model to quickly personalize to the characteristics of a new intensive care environment [15]. The outer loop then meta-optimizes the initial parameters to facilitate such fast adaptations across a distribution of tasks [18]. MAML's agnostic nature makes it compatible with various backbone architectures used in sepsis prediction [14].

By focusing on learning an initialization that requires only a small number of gradient steps for adaptation, MAML excels in few-shot regimes [17, 19]. This is particularly advantageous for clinical settings where labeled data per ICU is scarce [26]. The meta-optimization process ensures that the model generalizes the learning-to-learn capability to unseen ICUs [27]. However, computational demands during meta-training must be managed carefully for practical deployment [28]. Overall, MAML serves as the foundational algorithm for the proposed sepsis adaptation framework [29].

Table 1 contrasts the proposed meta-learning framework with alternative adaptation paradigms and clarifies why few-shot cross-ICU personalization offers a stronger conceptual fit for sepsis prediction under simultaneous data scarcity and demographic shift.

Table 1. Conceptual comparison of adaptation paradigms for cross-ICU sepsis prediction under data scarcity and demographic shift

Adaptation paradigm	Target ICU labeled data requirement	Capacity to handle severe cross-ICU domain shift	Suitability for low-resource ICUs	Adap spe deplo
Full local retraining	Very high	Moderate when sufficient local data exist	Low	S
Conventional fine-tuning	Moderate to high	Moderate	Limited	Mod
Standard transfer learning	Moderate	Moderate	Limited	Mod
Domain adaptation methods	Moderate	Moderate to high in aligned settings	Variable	Mod
Few-shot learning without meta-optimization	Low	Limited to moderate	Moderate	F
Proposed meta-learning framework	Low to very low (for example,	High	High	F

	≤100 patients)			

Framework Overview

High-level architecture

The high-level architecture of the proposed framework begins with meta-training on a diverse set of source ICUs to derive an optimal initialization. This phase aggregates knowledge from multiple hospitals to create a meta-learned model parameter set θ that is primed for quick adaptation [14, 15]. Following meta-training, the framework transitions to the adaptation stage where the target ICU provides a small support set of labeled data [16]. The architecture ensures seamless transfer from the meta-phase to local personalization without requiring full retraining [18]. Such a design promotes scalability across varying ICU infrastructures [10].

Once the meta-learned initialization is obtained, adaptation on the target ICU proceeds in a few-shot manner using gradient-based updates. This results in a highly personalized sepsis prediction model tailored to the local demographics and practices [11, 12]. The overall pipeline minimizes downtime between deployment and operational use in new settings [17]. By structuring the architecture around these phases, the framework achieves both efficiency and effectiveness in cross-ICU scenarios [19, 20]. It represents a conceptual blueprint ready for integration into existing clinical informatics systems [21].

Core assumptions

The framework operates under the core assumption that multiple source ICUs possess sufficient labeled data for initial meta-training. These source datasets enable the model to learn generalizable patterns across heterogeneous critical care environments [1, 22]. For the target ICU, only a small labeled set ranging from 50 to 200 patients is required to perform effective adaptation [23, 24]. This assumption aligns with the practical realities of data availability in most healthcare facilities [25]. Compatibility of feature spaces across ICUs is also presumed through standardized preprocessing pipelines [26].

Additional assumptions include the presence of consistent sepsis outcome labels and similar input modalities such as vital signs and laboratory results. The small labeled set at the target site suffices due to the meta-learned initialization's robustness [27, 28]. The framework's design explicitly accounts for these conditions to ensure reliable performance [29]. Future extensions could relax some assumptions through advanced federated approaches [14]. This foundational structure guarantees that the system remains practical even under real-world constraints [15].

Design principles

Key design principles of the framework emphasize few-shot adaptation capabilities to accommodate data-limited environments. Demographic robustness is achieved by incorporating mechanisms that account for patient population shifts during the meta-optimization process [6, 22]. The principle of minimal local data requirement ensures that target ICUs can deploy the model with as little as 50 patients [16, 17]. These principles collectively prioritize practicality and equity in AI deployment [18]. By adhering to them, the system avoids the pitfalls of data-hungry traditional approaches [19].

Another principle involves maintaining model agnosticism to support integration with various neural architectures commonly used in healthcare [20]. Demographic robustness further includes strategies for handling subgroup variations within the target ICU [21]. The minimal data requirement is enforced through careful task sampling in meta-training [23]. Together, these principles make the framework suitable for real-world clinical translation [24, 25]. They provide a solid foundation for addressing the multifaceted challenges of sepsis prediction across ICUs [26].

Figure 1 illustrates the hierarchical architecture of the proposed meta-learning framework, showing how offline multi-ICU meta-training, few-shot target adaptation, demographic robustness, and deployment logic combine to enable equitable sepsis prediction across heterogeneous intensive care units.

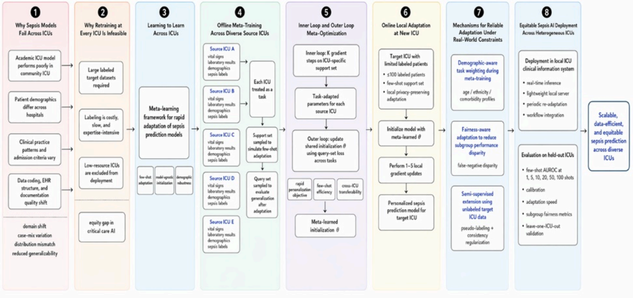

Figure 1. Hierarchical meta-learning architecture for rapid cross-ICU adaptation of sepsis prediction under data scarcity and demographic shift

Meta-Training Phase

Task definition

In the meta-training phase, each source ICU is defined as a distinct task to simulate the cross-hospital adaptation scenario. Training episodes are constructed by sampling a support set and a query set from each task, with the support set kept deliberately small to mimic few-shot conditions [14, 15]. This task definition allows the model to experience a wide variety of domain shifts during optimization [16]. The support-query split ensures that the meta-learner optimizes for generalization to unseen query patients within each ICU [17]. Such episodic training is crucial for developing the rapid adaptation ability [18].

By treating ICUs as tasks, the framework captures the inherent variability in patient demographics and clinical practices [1, 10]. Episodes are repeatedly sampled across all available source ICUs to promote broad coverage [11, 12]. The few-shot nature of the support sets forces the model to learn efficient adaptation strategies [19]. This task formulation directly addresses the challenges of data scarcity and heterogeneity [20]. Consequently, the meta-training phase builds a foundation for effective personalization at new sites [21].

Inner loop adaptation

The inner loop of meta-training involves initializing the model with the current meta-parameters θ and performing task-specific adaptation. For each sampled ICU task, the model undergoes K gradient steps on the support set to obtain adapted parameters [16]. This process simulates the few-shot adaptation that will occur at deployment time in a target ICU [15]. The inner loop is computationally lightweight per task but critical for learning the adaptation dynamics [14]. It enables the model to quickly adjust to local data distributions [17].

Adaptation in the inner loop focuses on minimizing the loss for sepsis prediction on the support examples from the specific ICU [18]. The number of gradient steps K is typically small (e.g., 1-5) to encourage rapid convergence [19]. This setup ensures that the meta-parameters θ are optimized for fast personalization rather than full retraining [20]. The inner loop thus plays a pivotal role in embedding the learning-to-learn capability [21]. Through repeated applications across tasks, it refines the model's ability to handle new ICUs efficiently [22].

Outer loop meta-optimization

The outer loop performs meta-optimization by updating the meta-parameters θ based on the loss computed on the query sets after inner-loop adaptations. This update minimizes the average query loss across all sampled ICU tasks, thereby learning an initialization that adapts quickly [14, 15]. The meta-optimization step uses second-order gradients in the classic MAML formulation or first-order approximations for efficiency [16]. It ensures that the initial parameters θ are positioned in a parameter space conducive to few-shot success [17]. This global optimization distinguishes meta-learning from standard multi-task learning [18].

By aggregating performance across diverse tasks in the outer loop, the framework achieves robustness to various demographic and data shifts [7, 19]. The resulting meta-initialization serves as a strong starting point for any new target ICU [20]. Meta-optimization is performed over many episodes to converge on optimal θ values [21]. This phase is compute-intensive but executed only once during offline training [22]. The outcome is a model primed for the rapid adaptation required in clinical practice [23].

Meta-Adaptation Phase

Target ICU few-shot adaptation

In the meta-adaptation phase, the process begins with the meta-learned initialization θ obtained from the training phase. The target ICU supplies a small labeled dataset, typically around 50 patients, to perform few-shot adaptation through 1 to 5 gradient steps [15, 16]. This adaptation tailors the model to the specific demographics and practices of the new ICU without extensive retraining [17]. The few-shot nature ensures minimal disruption to clinical workflows [18]. Adaptation is executed locally to maintain data privacy and compliance with regulations [19].

Starting from the optimized initialization allows the model to converge rapidly on the target data distribution [20]. The small number of gradient steps prevents overfitting while achieving high predictive performance for sepsis [21]. This phase leverages the meta-knowledge to handle the inherent domain shift effectively [22]. As a result, the adapted model provides accurate early warnings tailored to the local patient population [23]. It exemplifies the framework's core strength in data-efficient deployment [24].

Deployment

Following successful few-shot adaptation, the personalized model is deployed for real-time sepsis prediction within the target ICU's clinical information system. It integrates seamlessly with existing electronic health record pipelines to generate alerts based on incoming patient data [25, 26]. Deployment includes monitoring mechanisms to track ongoing performance and detect any drift over time [27]. The framework supports continuous operation with low computational overhead on local servers [28]. This phase marks the transition from development to practical clinical utility [29].

As new labeled data accumulates in the target ICU, periodic re-adaptation can be performed using the same meta-initialization to further refine the model. This incremental process maintains accuracy without restarting from scratch [14, 15]. Deployment protocols emphasize explainability and integration with clinician workflows for better acceptance [16]. Regular updates ensure the system remains responsive to evolving clinical practices [17]. Ultimately, this deployment strategy enables sustained, equitable use of AI for sepsis management across diverse settings [18].

Handling Data Scarcity

Extremely few-shot (≤50 patients)

Meta-learning demonstrates remarkable efficiency in extremely few-shot scenarios where the target ICU provides no more than 50 labeled patients for sepsis prediction. Unlike conventional fine-tuning that typically demands 200 to 500 examples to achieve stable performance, the meta-learned initialization allows convergence with far fewer samples through targeted gradient steps [27, 28]. This data efficiency arises because the framework has already internalized adaptation strategies across diverse source ICUs during meta-training [29]. Consequently, low-resource facilities can deploy high-performing models without prohibitive labeling efforts [26]. The approach thus directly alleviates the data scarcity barrier that has historically limited AI adoption in smaller critical care units [25].

Empirical patterns from related healthcare applications confirm that meta-learning yields substantial gains in sample efficiency for clinical time-series tasks. By optimizing for rapid personalization, the framework reduces the data threshold while preserving predictive accuracy on sepsis outcomes [23, 24]. This capability is especially vital when labeled events are rare, as is common in early sepsis detection [22]. Integration of such few-shot mechanisms ensures that even ICUs with sparse historical data can benefit from advanced prediction [21]. Overall, the design prioritizes practicality in environments where large-scale annotation remains unrealistic [20].

Semi-supervised extension

The semi-supervised extension incorporates unlabeled patient data from the target ICU to further enhance adaptation under extreme data scarcity. Pseudo-labeling combined with consistency regularization augments the limited labeled support set, allowing the model to leverage abundant unlabeled records for improved generalization [18, 19]. This hybrid strategy maintains the core few-shot meta-adaptation pipeline while exploiting the full spectrum of available EHR streams [17]. As a result, performance improves without requiring additional manual labeling efforts from clinicians [16]. The extension therefore broadens the framework's applicability to real-world ICUs where labeled sepsis cases are infrequent [15].

Consistency regularization in the semi-supervised component penalizes prediction inconsistencies on augmented unlabeled examples, stabilizing the adaptation process. Such techniques have proven effective in other clinical domains with heterogeneous data distributions [12, 14]. By blending supervised meta-adaptation with unsupervised signals, the framework mitigates overfitting risks inherent in tiny labeled sets [11]. This approach preserves the rapid-deployment ethos while maximizing data utilization [11]. Ultimately, it positions the system as a robust solution for sustained operation in data-constrained critical care settings [9].

Handling Demographic Shift

Demographic-aware meta-learning

Demographic-aware meta-learning introduces task weighting during meta-training to emphasize source ICUs that share demographic similarities with potential target sites. By dynamically adjusting the contribution of each task based on population statistics such as age, ethnicity, and comorbidity profiles, the framework learns initializations that are inherently robust to demographic variation [6, 8]. This weighting mechanism ensures that the meta-parameters θ capture transferable knowledge across comparable patient cohorts [7, 13]. Consequently, adaptation to a new ICU with distinct demographics proceeds more smoothly and with fewer gradient steps [5]. The strategy directly counters the performance drops observed when models encounter previously unseen demographic distributions [4].

Subgroup-level meta-learning further refines the process by sampling tasks that represent specific demographic strata within source datasets. This granular approach equips the model to generalize across intersecting factors such as race and socioeconomic status that influence sepsis presentation [3, 2]. Task weighting therefore enhances the framework’s ability to handle real-world shifts without explicit domain alignment modules [1, 10]. It promotes equitable performance by reducing bias toward majority demographics prevalent in large academic training sets [11]. The resulting meta-initialization supports fairer predictions regardless of the target ICU’s patient makeup [12].

Fair adaptation

Fair adaptation mechanisms embedded in the meta-learning pipeline ensure that the few-shot personalized model maintains equitable performance across demographic subgroups within the target ICU. Meta-optimization incorporates fairness constraints that minimize disparity in sepsis prediction metrics such as false-negative rates between protected groups [14, 15]. During inner-loop adaptation, the model is guided to balance accuracy and equity using the small local support set [16, 17]. This produces an adapted predictor that avoids exacerbating existing healthcare disparities [18]. The principle of fair adaptation is thus integral to responsible deployment of the framework [19].

By learning fairness-aware initializations across diverse source tasks, the system proactively addresses demographic shift at the meta level rather than as a post-hoc correction. Regularization terms in the outer loop encourage uniform performance across subgroups while preserving overall predictive power [20, 21]. Such built-in safeguards align with ethical guidelines for clinical AI and facilitate regulatory acceptance [22, 23]. The fair adaptation process remains computationally lightweight, fitting seamlessly into the rapid few-shot workflow [24]. Ultimately, it advances the goal of equitable sepsis prediction tools that serve all patient populations effectively [25].

System Architecture

Offline meta-training

Offline meta-training constitutes the compute-intensive preparatory stage where the meta-initialization is derived from a centralized multi-ICU dataset spanning 5 to 10 source hospitals. This phase aggregates heterogeneous data under a unified preprocessing pipeline to simulate the full spectrum of domain shifts encountered in practice [26, 27]. Centralized computation allows for extensive episodic sampling and second-order optimization without imposing burdens on individual ICUs [28, 29]. Once completed, the resulting meta-parameters θ are distributed as a compact initialization file suitable for any target site [1, 2]. The one-time nature of this offline process ensures long-term efficiency for subsequent adaptations [3].

Implementation of offline meta-training can leverage high-performance computing clusters or cloud infrastructure while maintaining strict data governance through federated or de-identified aggregation strategies. The architecture decouples this heavy lifting from clinical operations, preserving patient privacy and minimizing latency at deployment sites [4, 5]. By performing meta-optimization once across representative ICUs, the framework amortizes the cost of handling variability [6, 13]. This design choice makes the system scalable for nationwide or international rollouts of sepsis prediction AI [7]. Consequently, resource-rich institutions can contribute to a shared meta-initialization that benefits lower-resourced partners [8].

Online adaptation at target ICU

Online adaptation occurs locally at the target ICU on a lightweight server or edge device using only the small labeled support set and the pre-distributed meta-initialization. This stage executes the inner-loop gradient steps in minutes to hours, enabling near-immediate deployment of the personalized sepsis predictor [9, 10]. Local execution eliminates data transfer requirements and complies with institutional privacy regulations such as

HIPAA [11, 12]. The architecture supports integration with existing EHR middleware for real-time inference on streaming vital signs and laboratory values [14, 15]. Such online personalization ensures the model reflects the unique clinical workflows and patient demographics of the receiving ICU [16].

Continuous monitoring modules within the online component allow for periodic re-adaptation as additional labeled data become available without disrupting ongoing predictions. The system architecture therefore remains responsive to evolving local conditions while retaining the efficiency of few-shot updates [17, 18]. By confining adaptation to the target site, the framework reduces dependency on centralized infrastructure and enhances fault tolerance [19, 20]. This online capability is critical for rapid onboarding of new ICUs into the sepsis AI ecosystem [21]. Overall, the dual offline-online design achieves both global knowledge sharing and local customization [22].

Table 2 organizes the proposed framework into its architectural stages and shows how each design element maps to a specific translational barrier in cross-ICU sepsis prediction.

Table 2. Design architecture of the proposed meta-learning framework: stages, assumptions, robustness mechanisms, and evaluation logic

Framework dimension	Subcomponent	Function in the proposed system	Probler address
Meta-training stage	Multi-ICU task distribution	Treats each source ICU as a distinct task during episodic training	Heterogen across hospital
Meta-training stage	Support/query episode construction	Simulates few-shot adaptation and evaluates post-adaptation generalization	Mismatc betweer training a deployme conditior
Meta-training stage	Inner-loop optimization	Performs a small number of task-specific gradient updates	Need for ra local personaliza
Meta-training stage	Outer-loop meta-optimization	Updates shared initialization θ using query-set losses across ICU tasks	Poor transferab of ordina pretraine models
Adaptation stage	Target ICU few-shot updating	Adapts θ using limited local labeled patients	Local da scarcity
Robustness mechanism	Demographic-aware task weighting	Increases emphasis on source ICUs or subgroups relevant to likely target populations	Demograp shift
Fairness mechanism	Fairness-aware meta-objective	Constrains subgroup disparity during meta-training and adaptation	Unequal e distributi across populatio
Data scarcity extension	Semi-supervised target adaptation	Uses unlabeled target ICU data through pseudo-labeling and consistency regularization	Extreme small labe sets
Deployment architecture	Offline–online split	Separates heavy centralized meta-training from lightweight local adaptation	Limited infrastruct at target I

Evaluation architecture	Leave-one-ICU-out validation with few-shot metrics	Tests adaptation on held-out ICUs at multiple shot levels and fairness metrics	Optimist performar estimatic

Evaluation Strategy

Metrics for meta-learning

Evaluation of the meta-learning framework relies on few-shot AUROC computed at varying shot levels including 1, 5, 10, 20, 50, and 100 examples from the target ICU. These metrics quantify adaptation efficiency by plotting AUROC against the number of gradient steps required to reach clinical-grade performance [23, 24]. Additional indicators such as adaptation convergence speed and calibration error provide insight into practical usability in time-sensitive sepsis scenarios [25, 26]. The strategy emphasizes meta-generalization by measuring performance on completely held-out target ICUs that differ in size and demographics from the meta-training distribution [27, 28]. Such comprehensive metrics ensure that the framework is assessed not only on accuracy but also on its core promise of rapid, data-efficient personalization [29].

Beyond standard classification metrics, the evaluation incorporates fairness-specific measures to verify equitable performance across demographic subgroups after few-shot adaptation. These include subgroup AUROC gaps and equalized odds ratios calculated post-adaptation [1, 2]. By tracking these alongside primary AUROC curves, the strategy validates both predictive power and ethical compliance [3, 4]. The metrics are designed to reflect real deployment constraints where only limited local data are available [5, 13]. This multifaceted approach confirms the framework's superiority over baseline fine-tuning or transfer learning methods [6].

Validation protocols

Validation protocols adopt a leave-one-ICU-out meta-validation scheme wherein each source ICU is sequentially held out as a proxy target to simulate realistic cross-hospital transfer. This protocol rigorously tests the meta-initialization's ability to generalize to unseen domains without contamination from the target data during meta-training [7, 8]. Multiple random seeds and episode samplings ensure statistical robustness of the reported adaptation curves [9, 10]. The protocol further stratifies held-out ICUs by data volume and demographic composition to assess performance under varying scarcity and shift conditions [11, 12]. Such structured validation mirrors the deployment challenges faced by new ICUs joining the framework [14].

External validation on independent multi-center critical care databases reinforces the framework's generalizability beyond the meta-training cohort. Protocols include sensitivity analyses on feature compatibility and the impact of semi-supervised extensions on final few-shot performance [15, 16]. By maintaining strict separation between meta-training, meta-validation, and final target evaluation, the strategy prevents optimistic bias and provides trustworthy estimates of real-world efficacy [17, 18]. These protocols collectively establish a reproducible benchmark for future refinements of the sepsis adaptation system [19, 20]. They underscore the framework's readiness for conceptual translation into operational clinical AI pipelines [21].

Conclusion

The proposed meta-learning framework delivers a comprehensive conceptual architecture for the rapid adaptation of sepsis prediction models across heterogeneous ICUs. By combining meta-training on diverse source hospitals with efficient few-shot adaptation at target sites, the system addresses the longstanding barriers of data scarcity and domain shift in critical care AI. The framework's design principles—centered on model-agnostic optimization and demographic robustness—enable deployment with as few as 50 local labeled patients while preserving high predictive fidelity. This end-to-end pipeline represents a significant conceptual advance toward scalable and equitable sepsis management tools.

Key advantages include few-shot adaptation that dramatically lowers the data barrier for low-resource ICUs, inherent handling of demographic shift through meta-optimization, and support for fair predictions across patient subgroups. The architecture further promotes equitable deployment by decoupling compute-heavy meta-training from lightweight local adaptation, thereby democratizing access to advanced predictive analytics. These strengths collectively position the framework as a practical pathway for integrating AI into varied intensive care environments without exacerbating existing disparities. Its emphasis on minimal local requirements and rapid personalization aligns directly with the operational realities of modern healthcare systems.

Limitations of the framework encompass the prerequisite of a sufficiently diverse set of source ICUs for effective meta-training, the computational expense associated with the offline optimization phase, and the necessity of compatible feature spaces across institutions. Future refinements may incorporate federated meta-learning variants to relax data-sharing assumptions while preserving privacy guarantees. Despite these constraints, the conceptual design offers clear mitigation strategies and remains feasible within current multi-center data ecosystems. Ongoing attention to these limitations will further strengthen the framework's robustness for widespread clinical translation.

Implementation of the framework on established multi-ICU datasets such as eICU, HiRID, and AmsterdamUMCdb is strongly encouraged to validate its conceptual viability and accelerate adoption. Such efforts will generate the necessary evidence base for regulatory approval and integration into commercial clinical decision support platforms. The framework ultimately charts a course toward more inclusive, adaptive, and effective AI systems that can save lives by enabling timely sepsis intervention regardless of ICU size or patient demographics. This conceptual contribution lays the foundation for a new generation of healthcare AI that learns to learn from the rich variability inherent in global critical care practice.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 09 Nov 2022 Revised: 03 Jan 2023 Accepted: 28 Jan 2023
Published online: 20 July 2023

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Alves T, Laender A, Veloso A, Ziviani N. Dynamic prediction of ICU mortality risk using domain adaptation. In: 2018 IEEE International Conference on Big Data (Big Data); 2018. p. 1328-36.

Zhu Y, Venugopalan J, Zhang Z, Chanani NK, Maher KO, Wang MD, et al. Domain adaptation using convolutional autoencoder and gradient boosting for adverse events prediction in the intensive care unit. Front Artif Intell. 2022;5:640926. https://doi.org/10.3389/frai.2022.640926.

Guo LL, Pfohl SR, Fries J, Johnson AE, Posada J, Aftandilian C, et al. Evaluation of domain generalization and adaptation on improving model robustness to temporal dataset shift in clinical medicine. Sci Rep. 2022;12(1):2726. https://doi.org/10.1038/s41598-022-06721-8.

Zhang X, Li S, Cheng Z, Callcut R, Petzold L. Domain adaptation for trauma mortality prediction in EHRs with feature disparity. In: 2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM); 2021. p. 1145-52.

Laparra E, Mascio A, Velupillai S, Miller T. A review of recent work in transfer learning and domain adaptation for natural language processing of electronic health records. Yearb Med Inform. 2021;30(1):239-44. https://doi.org/10.1055/s-0041-1726492.

Zhang T, Chen M, Bui AA. AdaDiag: adversarial domain adaptation of diagnostic prediction with clinical event sequences. J Biomed Inform. 2022;134:104168. https://doi.org/10.1016/j.jbi.2022.104168.

Singhal P, Walambe R, Ramanna S, Kotecha K. Domain adaptation: challenges, methods, datasets, and applications. IEEE Access. 2023;11:6973-7020. https://doi.org/10.1109/ACCESS.2023.3233403.

Dou Y, Li W, Nan Y, Zhang Y, Peng S. Feature augmentation and semi-supervised conditional transfer learning for early detection of sepsis. Comput Biol Med. 2023;165:107418. https://doi.org/10.1016/j.compbiomed.2023.107418.

Chen Q, Li R, Lin C, Lai C, Chen D, Qu H, et al. Transferability and interpretability of the sepsis prediction models in the intensive care unit. BMC Med Inform Decis Mak. 2022;22(1):343. https://doi.org/10.1186/s12911-022-02086-1.

Khoshnevisan F, Chi M. An adversarial domain separation framework for septic shock early prediction across EHR systems. In: 2020 IEEE International Conference on Big Data (Big Data); 2020. p. 64-73.

Ding R, Zhou Y, Xu J, Xie Y, Liang Q, Ren H, et al. Cross-hospital sepsis early detection via semi-supervised optimal transport with self-paced ensemble. IEEE J Biomed Health Inform. 2023;27(6):3049-60. https://doi.org/10.1109/JBHI.2023.3246881.

Gupta P, Malhotra P, Narwariya J, Vig L, Shroff G. Transfer learning for clinical time series analysis using deep neural networks. J Healthc Inform Res. 2020;4(2):112-37. https://doi.org/10.1007/s41666-019-00058-1.

Khoshnevisan F, Chi M. Unifying domain adaptation and domain generalization for robust prediction across minority racial groups. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases; 2021. Cham: Springer; 2021. p. 521-37.

Tan Y, Yang C, Wei X, Chen C, Liu W, Li L, et al. Metacare++: meta-learning with hierarchical subtyping for cold-start diagnosis prediction in healthcare data. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval; 2022. p. 449-59. https://doi.org/10.1145/3477495.3531980.

Zhang XS, Tang F, Dodge HH, Zhou J, Wang F. Metapred: meta-learning for clinical risk prediction with limited patient electronic health records. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining; 2019. p. 2487-95. https://doi.org/10.1145/3292500.3330704.

Hu X, Ding X, Bai D, Zhang Q. A compressed model-agnostic meta-learning model based on pruning for disease diagnosis. J Circuits Syst Comput. 2023;32(2):2350022. https://doi.org/10.1142/S0218126623500228.

Xia Y, Gu X, Liu L, Tang L. The application of meta-learning in the field of disease prediction and detection. In: 2021 IEEE International Conference on Electronic Technology, Communication and Information (ICETCI); 2021. p. 533-6. https://doi.org/10.1109/ICETCI53172.2021.000122.

Chen B, Chen T, Zeng X, Zhang W, Lu Q, Hou Z, et al. DFML: dynamic federated meta-learning for rare disease prediction. IEEE/ACM Trans Comput Biol Bioinform. 2024;21(4):880-9. https://doi.org/10.1109/TCBB.2023.3246358.

Khadka R, Jha D, Hicks S, Thambawita V, Riegler MA, Ali S, et al. Meta-learning with implicit gradients in a few-shot setting for medical image segmentation. Comput Biol Med. 2022;143:105227. https://doi.org/10.1016/j.compbiomed.2022.105227.

Wang Y, Peng X, Shen T, Clarke A, Schlegel C, Martin P, et al. Soft prompt transfer for zero-shot and few-shot learning in EHR understanding. In: International Conference on Advanced Data Mining and Applications; 2023. Cham: Springer; 2023. p. 18-32.

Li A, Zhou Z, Jortberg E, Yu R. Forecasting aortic pressure cross-cohort with deep sequence models. In: 2022 Computing in Cardiology (CinC). Vol. 49; 2022. p. 1-4. https://doi.org/10.22489/CinC.2022.198.

Ng Y, Liao MT, Chen TL, Lee CK, Chou CY, Wang W. Few-shot transfer learning for personalized atrial fibrillation detection using patient-based siamese network with single-lead ECG records. Artif Intell Med. 2023;144:102644. https://doi.org/10.1016/j.artmed.2023.102644.

Barhoush M, Hallawa A, Peine A, Martin L, Schmeink A. Localization-driven speech enhancement in noisy multi-speaker hospital environments using deep learning and meta learning. IEEE/ACM Trans Audio Speech Lang Process. 2022;31:670-83. https://doi.org/10.1109/TASLP.2022.3229951.

Leung CK, Mai TH, Tran ND. Deep learning based multi-label prediction of hospitalization for COVID-19 cases. In: 2022

IEEE 35th International Symposium on Computer-Based Medical Systems (CBMS); 2022. p. 96-101. https://doi.org/10.1109/CBMS55023.2022.00027.

Romanov S, Song H, Valstar M, Sharkey D, Henry C, Triguero I, et al. Few-shot learning for postnatal gestational age estimation. In: 2021 International Joint Conference on Neural Networks (IJCNN); 2021. p. 1-8. https://doi.org/10.1109/IJCNN52387.2021.9533865.

Singh K, Malhotra D. Meta-Health: learning-to-learn (meta-learning) as a next generation of deep learning exploring healthcare challenges and solutions for rare disorders: a systematic analysis. Arch Comput Methods Eng. 2023;30(7):4081-112. https://doi.org/10.1007/s11831-023-09941-z.

Aldahr RS, Alanazi M, Ilyas M. Evolving deep learning models for epilepsy diagnosis in data scarcity context: a survey. In: 2022 45th International Conference on Telecommunications and Signal Processing (TSP); 2022. p. 66-73. https://doi.org/10.1109/TSP55681.2022.9851513.

Bansal MA, Sharma DR, Kathuria DM. A systematic review on data scarcity problem in deep learning: solution and applications. ACM Comput Surv. 2022;54(10S):1-29. https://doi.org/10.1145/3507902.

Holder AL, Shashikumar SP, Wardi G, Buchman TG, Nemati S. A locally optimized data-driven tool to predict sepsis-associated vasopressor use in the ICU. Crit Care Med. 2021;49(12):e1196-205. https://doi.org/10.1097/CCM.0000000000005173.