ORIGINAL RESEARCH Open access

Hierarchical Transformer Model with Uncertainty-Aware Attention for Predicting Delayed Cerebral Ischemia after Aneurysmal Subarachnoid Hemorrhage Using Hourly Neurological and Vital Sign Data

John Miller[1*], Mark Davis[1], Ethan Clark[2]

Abstract

Delayed cerebral ischemia (DCI) following aneurysmal subarachnoid hemorrhage is a significant cause of morbidity, mortality, and long-term neurological disability. Current clinical scores like WFNS, Hunt-Hess, and modified Fisher scale provide useful baseline risk information but often fail to capture subtle multi-hour deteriorations. Standard recurrent models can process sequential data but struggle with long-term dependencies and do not offer clinicians useful uncertainty information. To address this, a hierarchical Transformer model is proposed for DCI prediction, leveraging short-term hourly changes and longer multi-day trends in neurological and vital sign data. The model incorporates an uncertainty-aware attention mechanism to minimize the impact of unreliable or missing data and generates risk-stratified alerts with confidence levels. This approach aims to offer an explainable, clinically actionable tool that supports early recognition of DCI while ensuring clinician oversight. Future work will involve retrospective development and prospective validation to enhance its clinical utility.

Keywords Clinical time series, Hierarchical transformer, Delayed cerebral ischemia, Aneurysmal subarachnoid hemorrhage, Uncertainty-aware attention, Neurocritical care

*Correspondence:
John Miller
john.miller@gmail.com
[1] Department of Healthcare AI Engineering, University of California, Berkeley, United States
[2] Department of Intelligent Clinical Systems, Stanford University, Stanford, United States

Introduction

Delayed cerebral ischemia is a clinically consequential complication after aneurysmal subarachnoid hemorrhage because it may develop after aneurysm securing and contribute to secondary brain injury. Contemporary aneurysmal subarachnoid hemorrhage guidelines emphasize early recognition of neurological deterioration, structured monitoring, and timely rescue therapy when DCI or vasospasm is suspected [1, 2]. Diagnostic consensus work further frames DCI as a syndrome that requires integration of clinical examination, imaging, hemodynamic context, and exclusion of alternative causes of decline [3]. A prediction framework for DCI should therefore model evolving neurological and physiological trajectories rather than rely only on static admission severity.

Hourly neurological assessment and vital-sign monitoring create a dense stream of clinically meaningful information in the neuro-intensive care unit. Glasgow Coma Scale trends, pupillary findings, new motor deficits, mean arterial pressure, oxygen saturation, intracranial pressure, and cerebral perfusion pressure may each carry partial information about impending deterioration when interpreted over time [2, 3]. Machine-learning studies in aSAH have shown that structured clinical variables can support DCI prediction, but many models remain closer to static or low-frequency risk stratification than to continuous temporal reasoning [4, 5]. Longitudinal modeling work suggests that changing risk profiles after aSAH are clinically important, making hourly data especially relevant for a conceptual DCI prediction framework [6].

Standard recurrent neural networks and early attention-based clinical models provided important foundations for sequential healthcare prediction, but they were not designed primarily for multi-day, irregular, uncertainty-rich neurocritical care data. Che, Purushotham, Cho, Sontag and Liu demonstrated how recurrent networks could account for missing values in multivariate medical time series, while Choi, Bahadori, Kulas, Schuetz, Stewart and Sun introduced reverse-time attention for interpretable healthcare prediction [7, 8]. TERTIAN and self-supervised clinical Transformers extend temporal representation by using attention mechanisms for sparse or irregular clinical data, while temporal fusion architectures show how multi-horizon forecasting can combine static and dynamic covariates [9-11]. However, DCI prediction requires not only long-range temporal modeling but also explicit handling of uncertainty from missing observations, charting artifacts, and ambiguous neurological examinations [12].

This article proposes a conceptual framework for a hierarchical Transformer with uncertainty-aware attention to predict delayed cerebral ischemia after aneurysmal subarachnoid hemorrhage using hourly neurological and vital-sign data. The framework integrates clinical DCI monitoring principles with advances in Transformer-based time-series modeling, uncertainty-aware prediction, and interpretable healthcare artificial intelligence [1, 9, 12]. Its central argument is that DCI prediction should combine local hourly trend detection, cross-day context, missingness-aware representation, and prediction-level uncertainty. The following sections define the clinical and computational background, describe the model architecture, and explain how uncertainty-aware attention could support safer neurocritical care decision support.

Background

Pathophysiology of DCI

Delayed cerebral ischemia after aneurysmal subarachnoid hemorrhage is typically understood as a delayed secondary injury process involving large-vessel vasospasm, impaired autoregulation, microthrombosis, cortical spreading depolarization, inflammation, and perfusion failure. The vulnerable period often begins several days after ictus and extends through the second week, which makes temporal surveillance central to clinical management [1, 3]. Radiological scale reviews show that early blood burden and hemorrhage distribution remain important predictors, but these baseline features do not fully capture evolving physiological risk during the ICU stay [13]. A time-series framework should therefore treat DCI as a dynamic process rather than as a fixed event determined only at admission.

Clinical monitoring in aSAH

Clinical monitoring in aSAH depends on repeated neurological examinations, hemodynamic assessment, and integration of bedside variables that may change hour by hour. Neurocritical care guidance emphasizes neurological examination, vital-sign control, intracranial pressure management when indicated, and surveillance for clinical deterioration or vasospasm [2]. In practice, charted measurements such as Glasgow Coma Scale, limb strength, cranial nerve findings, mean arterial pressure, heart rate, intracranial pressure, and cerebral perfusion pressure create an event stream that is both clinically rich and irregularly complete. This framework treats those observations as temporally ordered tokens whose meaning depends on both their value and their timing [3, 6].

Existing DCI prediction models

Existing DCI prediction approaches include traditional grading systems, radiological scales, regression models, random forests, and broader machine-learning pipelines. Ramos and colleagues showed that machine learning could improve DCI prediction after subarachnoid hemorrhage, while Savarraj and colleagues developed machine-learning models for DCI and outcome prediction using structured clinical information [4, 5]. Recent systematic reviews and meta-analyses have synthesized growing evidence for machine-learning-based DCI prediction but also highlight heterogeneity in data sources, predictors, validation strategies, and model transparency [14-16]. These limitations support a conceptual shift toward models that represent longitudinal bedside data, report uncertainty, and remain interpretable to clinicians.

Transformer architectures for clinical time series

Transformer architectures are attractive for clinical time series because self-attention can connect observations across long temporal distances without relying solely on recurrent state propagation. TERTIAN introduced a time-aware Transformer-based hierarchical attention network for ICU endpoint prediction, and self-supervised clinical Transformers have further shown how sparse and irregular electronic health record sequences can be represented with attention mechanisms [9, 10]. Multi-time attention networks and temporal fusion Transformers add complementary ideas for handling irregular sampling, covariate interactions, and multi-horizon temporal forecasting [11, 17]. For DCI prediction, the conceptual value of a hierarchical Transformer is that hourly changes can be modeled locally while cross-day patterns remain available to later attention layers.

Uncertainty-aware attention

Uncertainty-aware attention refers to mechanisms that estimate uncertainty at the token, feature, time-window, or sequence level and use that estimate to modify contextualization. Healthcare time-series work on uncertainty-aware self-attention, uncertainty-aware disease prediction, and uncertainty-aware deep ensembles suggests that uncertainty can be embedded into prediction pipelines rather than appended only after model training [12, 18, 19]. In DCI monitoring, this matters because a missing examination, a physiologically implausible vital sign, or a transient artifact should not be treated as equally reliable evidence. A credal or evidential attention design would allow the framework to represent not only what the model attends to, but also how trustworthy each attended observation appears.

Framework Overview

High-level architecture

The proposed architecture would ingest hourly neurological assessments and vital signs as multivariate temporal tokens and pass them through time-aware positional encoding before hierarchical Transformer processing. Local attention blocks would summarize short windows of observations, while higher-level blocks would relate these summaries across several days of the DCI-risk period [9, 11]. Uncertainty-aware attention would modulate contextualization so that missing, noisy, or clinically inconsistent observations receive lower effective influence when appropriate [12, 18]. The final prediction head would output a DCI probability together with a prediction-level uncertainty interval suitable for clinical decision support.

Figure 1 presents the proposed hierarchical Transformer architecture, showing how hourly neurological and vital-sign observations are transformed into uncertainty-calibrated, interpretable, clinician-supervised delayed cerebral ischemia risk support.

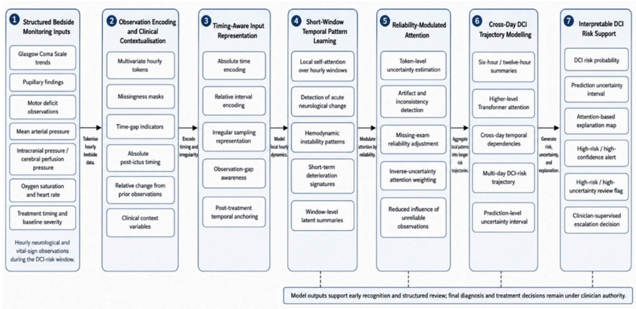

Figure 1. Hierarchical Transformer architecture with uncertainty-aware attention for delayed cerebral ischemia prediction after aneurysmal subarachnoid hemorrhage

Core assumptions

The framework assumes that structured hourly data are available for the first 72 to 120 hours after aneurysm treatment, while allowing later extension across the broader DCI surveillance window. This assumption is clinically motivated by guideline-based neuro-ICU monitoring and by dynamic DCI risk studies showing that risk changes over time rather than remaining fixed after admission [2, 6]. Minimal preprocessing would encode observed values, missingness masks, time gaps, and clinically relevant context such as aneurysm treatment timing or baseline severity. Rather than requiring perfect data completeness, the model would treat irregularity and missingness as part of the clinical signal.

Design principles

The framework is guided by three design principles: uncertainty calibratability, interpretability through hierarchical attention, and avoidance of interpolation artifacts. Time-series imputation models such as SAITS show that self-attention can learn complex missingness structures, while recurrent missingness-aware models demonstrate that missing values themselves may contain predictive information [7, 20]. However, for neurocritical care decision support, aggressive interpolation could create artificial stability or hide clinically meaningful observation gaps. The proposed design therefore represents missingness explicitly and allows attention to operate over observed events, masks, and time differences.

Table 1 clarifies how the proposed architecture translates specific neurocritical-care monitoring challenges into corresponding computational design choices.

Table 1. Conceptual Mapping between Clinical DCI Monitoring Needs and Architectural Design Responses

Clinical monitoring need in aSAH	Limitation of conventional/static prediction	Proposed architectural response	Ac cond valu D pred
DCI risk evolves across hours and days	Admission scores may not capture delayed deterioration	Hierarchical temporal modelling across hourly and multi-day windows	Repr DC dyn traj rathe fi base
Neurological decline may be subtle or intermittent	Simple sequential models may miss weak multi-hour patterns	Local hourly Transformer blocks	De short- neur physi deter
Risk depends on longer post-ictus context	Recurrent models may struggle with long-range dependencies	Higher-level cross-day Transformer aggregation	Links char broad risk
ICU data are irregular and incomplete	Forced interpolation may create artificial stability	Explicit missingness masks and time-gap encoding	Pres the mea abs de obse
Charted values may include artifacts or unreliable entries	Standard attention treats attended tokens as equally credible	Uncertainty-aware attention gating	Re influ n incor or m obse
Clinicians need usable risk information	Binary alerts may obscure confidence and ambiguity	Risk score plus uncertainty interval	Distin cor wa from req fu re
Bedside adoption requires interpretability	Black-box predictions may reduce trust	Attention-based explanation by variable and time window	Show obse and t seg shap a

Hierarchical Multi-Scale Transformer

Time-aware positional encoding

Time-aware positional encoding would represent both absolute timestamps and relative time gaps between observations so that the model can distinguish a stable hourly sequence from a sparse or interrupted sequence. This is necessary because ICU and electronic health record data are often irregular even when monitoring is expected to be frequent [10, 17]. Time-aware Transformer designs and multi-time attention networks provide conceptual support for encoding timing directly rather than assuming evenly spaced measurements [9, 17]. In DCI prediction, this would allow the model to interpret a neurological change differently depending on whether it follows a recent stable examination or a long gap in charted assessments.

Local-global attention

Local-global attention would combine short-window self-attention for immediate trend detection with cross-day attention for longer DCI-risk trajectories. A local block might emphasize a sudden fall in mean arterial pressure, new focal weakness, or worsening level of consciousness, while a global block might connect this change to prior blood burden, earlier instability, or persistent vasospasm concern [1, 3, 13]. Temporal fusion and multi-horizon forecasting models illustrate how short-term signals and longer forecasting contexts can be combined within attention-based architectures [11, 21]. This design is especially appropriate for DCI because the clinically relevant pattern may unfold over hours while being conditioned by events across several days.

Hierarchical aggregation

Hierarchical aggregation would summarize attention outputs over clinically meaningful intervals, such as six-hour or twelve-hour blocks, before passing them to higher-level Transformer layers and the final prediction head. This structure follows the logic of hierarchical clinical attention models in which lower layers encode local temporal information and higher layers represent broader clinical context [9]. For bedside interpretability, aggregation by clinically familiar time windows could help clinicians see whether risk was driven by overnight neurological decline, sustained hemodynamic instability, or later-phase deterioration. The model would therefore aim to reduce noise without hiding the temporal origin of the DCI warning.

Uncertainty-Aware Attention Mechanism

Uncertainty estimation

The uncertainty module would estimate token-level epistemic uncertainty for each hourly observation before or during attention computation. Uncertainty-aware self-attention models and uncertainty-aware deep ensembles suggest that reliability information can be derived from model disagreement, dropout-based approximations, or evidential representations [12, 19]. In this framework, a single neurological assessment with inconsistent surrounding vital signs, an implausible physiological value, or a missing examination token could receive a higher uncertainty estimate. That estimate would then become part of the attention mechanism rather than only a post-hoc explanation.

Attention score modulation

Attention score modulation would multiply or otherwise transform raw attention weights using an inverse-uncertainty factor so that unreliable observations contribute less to contextual representation. Time interval uncertainty-aware disease prediction work supports the idea that uncertainty in timing and observation quality can be incorporated into longitudinal healthcare prediction [18]. In DCI prediction, this mechanism could reduce the influence of artifacts, missing values, or isolated abnormalities that lack corroborating neurological or physiological context. The resulting attention distribution would represent both temporal relevance and estimated reliability.

Output-level uncertainty

Output-level uncertainty would be generated by an ensemble of moderately sized hierarchical Transformers or by Monte Carlo dropout applied to the prediction head and selected attention layers. Uncertainty-aware ensemble methods for clinical time series show how model uncertainty can support more reliable and explainable prediction when clinical data are noisy or incomplete [19]. In the proposed DCI framework, the prediction would therefore include both a risk estimate and an uncertainty interval rather than a single deterministic score. This approach could preserve interpretability while helping clinicians distinguish confident alerts from cases where additional neurological assessment, imaging, or expert review may be needed.

Handling Irregularly Sampled Data

Missing data as a feature

Missing data in neurocritical care should not be treated only as a technical nuisance, because the absence of a neurological examination or vital-sign entry may reflect workflow disruption, clinical instability, sedation, imaging transfer, or emergency intervention. Missingness-aware recurrent models and self-attention imputation frameworks show that masks and observation patterns can contain predictive structure in multivariate medical time series [7, 20]. Multi-headed Transformer models for charted vital signs further support the idea that irregularly available physiological measurements can be represented as structured temporal inputs rather than discarded or homogenized [22]. In the proposed framework, each hourly token would therefore include observed values, missingness indicators, and timing information so that the model could learn when absence of data is itself clinically meaningful.

Interpolation-free processing

Interpolation-free processing would avoid forcing hourly neuro-ICU data into artificially complete trajectories through forward filling or linear interpolation. Self-supervised and irregularly sampled clinical Transformers demonstrate that attention-based models can operate over sparse medical sequences when time gaps are encoded directly [10, 23]. ICU vital-sign forecasting studies also show that high-frequency physiological trajectories can be modeled as multivariate temporal processes without reducing them to static summaries [21, 24]. For DCI prediction, this approach would allow the framework to distinguish true clinical stability from a charting gap that merely appears stable after imputation.

Clinical Interpretability

Attention visualisation

The interpretability module would generate attention visualisations that identify which hourly neurological observations, physiological signals, and temporal windows most influenced the DCI risk estimate. RETAIN and Dipole established that attention mechanisms can support clinically interpretable prediction by linking model outputs to specific visits, variables, or temporal contexts [8, 25]. In this framework, attention heatmaps could highlight an overnight reduction in Glasgow Coma Scale, a cluster of new motor deficits, or a sustained pattern of hemodynamic instability that contributed to the warning. Hierarchical visualisation would also allow clinicians to distinguish local short-window drivers from broader multi-day risk patterns.

Explanation of uncertainty

The explanation layer would not only show what the model attended to but also explain when the model was uncertain about its own prediction. Uncertainty-aware deep ensembles and uncertainty-aware self-attention models suggest that predictive reliability can be communicated alongside feature relevance, which is essential when clinical data are sparse, noisy, or contradictory [12, 19]. In a neuro-ICU workflow, a high-risk but high-uncertainty alert might recommend additional neurological checks, review of charted vital signs, or confirmatory imaging before escalation. This design would make uncertainty a visible part of clinical reasoning rather than an invisible property of the algorithm.

Evaluation Strategy

Discrimination and calibration

A future retrospective evaluation would conceptually assess discrimination using AUROC and AUPRC while avoiding claims of performance before empirical validation. Calibration would be equally important because DCI alerts must align predicted risk with observed risk if clinicians are to use them safely for escalation decisions. Existing machine-learning studies and systematic reviews of DCI prediction show that model comparisons should include clinically meaningful baselines such as logistic regression, random forest, grading scales, and previously proposed machine-learning pipelines [4, 14-16]. The hierarchical Transformer would therefore be evaluated conceptually against recurrent, Transformer-based, and conventional clinical prediction approaches rather than presented as empirically superior.

Uncertainty-aware metrics

Uncertainty-aware evaluation would examine whether the framework's confidence estimates are clinically coherent, not merely whether the risk score separates future DCI from non-DCI cases. Brier score, credible interval coverage, uncertainty-calibration slope, and deferral behavior would help assess whether high uncertainty corresponds to ambiguous or insufficient data conditions [12, 19]. Recent DCI prediction models using machine learning, inflammatory markers, nomograms, and early grading scales illustrate the diversity of candidate predictors, but most do not make uncertainty a central output for clinical workflow design [26-30]. A conceptual uncertainty-aware evaluation would therefore ask whether the model knows when its own alert should be treated cautiously.

Clinical utility simulation

Clinical utility simulation would examine how the framework might behave under a policy such as alerting only when predicted risk exceeds a threshold and prediction uncertainty remains below a threshold. This differs from risk-only alerting because it allows the system to defer uncertain cases rather than generate a confident-looking warning from incomplete or unreliable data. Transcranial Doppler biomarker reviews and radiological prediction studies show that DCI assessment already depends on combining imperfect signals, which makes uncertainty-sensitive alerting conceptually aligned with real clinical decision-making [13, 31]. The simulation would not claim outcome improvement but would explore how uncertainty-aware triage might reduce unnecessary escalation while preserving attention to high-risk patterns.

Limitations

Technical limitations

The proposed framework would depend on structured hourly data that may not be consistently available across all neuro-intensive care units. Extending the receptive field across the full DCI vulnerability window could impose substantial memory and computational demands, especially when multivariate neurological, hemodynamic, and treatment variables are modeled at fine temporal resolution [9, 17, 32]. Uncertainty modules, ensembles, and Monte Carlo dropout would add further computational cost and may complicate real-time deployment [12, 19]. These limitations mean that the framework should be viewed as a design proposal requiring pragmatic adaptation to local data quality and infrastructure.

Clinical limitations

The framework could not replace expert neurological assessment, continuous invasive neuromonitoring, transcranial Doppler surveillance, electroencephalography, perfusion imaging, or clinician judgment in complex aSAH cases. Guidelines and consensus recommendations emphasize that DCI diagnosis requires integration of clinical deterioration, imaging, hemodynamics, and exclusion of alternative causes, which no single model can fully automate [1-3]. Poor calibration or excessive alert frequency could worsen alarm fatigue, especially if uncertainty estimates are not displayed clearly or if clinicians distrust the rationale behind alerts. Prospective validation would therefore be essential before the framework could be considered for clinical deployment.

Impact on Neurocritical Care Workflow

Integration into EMR

The proposed model could run as a decision-support service within the neuro-ICU electronic medical record, automatically ingesting hourly neurological assessments, vital signs, missingness masks, and treatment context. Transformer-based clinical time-series models and ICU forecasting frameworks support the feasibility of representing continuously updated structured data streams for prediction tasks [21-32]. In practice, the framework would need to present risk, uncertainty, and explanation together so that clinicians can interpret alerts without leaving their usual workflow. Integration should therefore prioritize transparent display and low-friction review rather than isolated algorithmic output.

Clinical decision support

Clinical decision support would distinguish high-confidence DCI risk alerts from low-confidence warnings that require further assessment before action. The framework could label uncertain predictions as low-confidence outputs and suggest review of neurological examinations, vital-sign artifacts, or missing charted data before recommending confirmatory imaging. When both risk and confidence are high, clinicians might consider closer examination, vascular imaging, perfusion assessment, or rescue-therapy discussion in accordance with institutional practice and guideline-based care [1-3]. This approach would position the model as a structured warning system rather than as an autonomous diagnostic authority.

Table 2 proposes a risk–uncertainty interpretation logic that positions the model as a clinician-supervised warning system rather than an autonomous diagnostic tool.

Table 2. Risk–Uncertainty Decision-Support Logic for Clinician-Supervised DCI Alert Interpretation

Predicted DCI risk	Prediction uncertainty	Suggested alert interpretation	Clinician-facing action logic	S
Low risk	Low uncertainty	Reassuring low-risk signal	Continue routine monitoring according to institutional protocol	e
Low risk	High uncertainty	Insufficient confidence in low-risk estimate	Review missing data, recent examination quality, and monitoring gaps	fi
Moderate risk	Low uncertainty	Stable intermediate concern	Increase clinical vigilance and compare	S i

			with recent neurological trajectory	
Moderate risk	High uncertainty	Ambiguous intermediate warning	Repeat neurological assessment or verify questionable vital-sign entries	u
High risk	Low uncertainty	High-priority DCI concern	Prompt clinician review, consider imaging or rescue-therapy discussion according to guidelines and local practice	a
High risk	High uncertainty	High-risk but unreliable signal	Escalate for expert review while checking artifacts, sedation, charting gaps, and alternative causes	A p
Rapidly rising risk	Any uncertainty level	Temporal deterioration pattern	Compare attention map with clinical events over the preceding hours	c
Discordant risk and clinical impression	Any uncertainty level	Model–clinician mismatch	Require human adjudication; do not treat model output as diagnostic authority	

Conclusion

The proposed conceptual framework combines a hierarchical Transformer with uncertainty-aware attention to support prediction of delayed cerebral ischemia after aneurysmal subarachnoid hemorrhage. It is designed around hourly neurological assessments and vital-sign data because these observations represent the evolving bedside state of the patient. The framework would model both short-term deterioration and longer multi-day risk trajectories during the vulnerable post-ictus period.

Its key advantage is the combination of multi-scale temporal representation, explicit uncertainty estimation, and interpretable attention mapping. Local attention would identify immediate clinical changes, while higher-level attention would connect these changes to broader temporal patterns. Uncertainty-aware attention would help prevent unreliable observations from being treated as equally trustworthy evidence.

The framework remains data-dependent and should be considered a conceptual design rather than a validated clinical tool. Its reliability would depend on the completeness, consistency, and clinical meaning of structured neuro-ICU data. Prospective validation would be necessary to determine whether it improves recognition of delayed cerebral ischemia without increasing alarm fatigue or unnecessary interventions.

Future work should implement the framework on multi-centre aneurysmal subarachnoid hemorrhage datasets and evaluate it in clinically realistic settings. Pilot studies should examine how clinicians interpret the model’s risk estimates, uncertainty intervals, and attention-based explanations. The ultimate goal would be to determine whether uncertainty-calibrated decision support can improve the timing, confidence, and safety of DCI recognition in neurocritical care.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 23 Jan 2026 Revised: 25 Jan 2026 Accepted: 20 Mar 2026
Published online: 20 July 2026

References

Hoh BL, Ko NU, Amin-Hanjani S, Chou SH, Cruz-Flores S, Dangayach NS, et al. 2023 guideline for the management of patients with aneurysmal subarachnoid hemorrhage: a guideline from the American Heart Association/American Stroke Association. Stroke. 2023;54(7):e314-70. https://doi.org/10.1161/STR.0000000000000436.

Treggiari MM, Rabinstein AA, Busl KM, Caylor MM, Citerio G, Deem S, et al. Guidelines for the neurocritical care management of aneurysmal subarachnoid hemorrhage. Neurocrit Care. 2023;39(1):1-28. https://doi.org/10.1007/s12028-023-01713-8.

Abdulazim A, Heilig M, Rinkel G, Etminan N. Diagnosis of delayed cerebral ischemia in patients with aneurysmal subarachnoid hemorrhage and triggers for intervention. Neurocrit Care. 2023;39(2):311-9. https://doi.org/10.1007/s12028-023-01741-4.

Ramos LA, van der Steen WE, Barros RS, Majoie CB, van den Berg R, Verbaan D, et al. Machine learning improves prediction of delayed cerebral ischemia in patients with subarachnoid hemorrhage. J Neurointerv Surg. 2019;11(5):497-502. https://doi.org/10.1136/neurintsurg-2018-014305.

Savarraj JP, Hergenroeder GW, Zhu L, Chang T, Park S, Megjhani M, et al. Machine learning to predict delayed cerebral ischemia and outcomes in subarachnoid hemorrhage. Neurology. 2021;96(4):e553-62. https://doi.org/10.1212/WNL.0000000000011242.

Willms JF, Inauen C, Bögli SY, Muroi C, Boss JM, Keller E. A longitudinal model for a dynamic risk score to predict delayed cerebral ischemia after subarachnoid hemorrhage. Bioengineering (Basel). 2024;11(10):988. https://doi.org/10.3390/bioengineering11100988.

Che Z, Purushotham S, Cho K, Sontag D, Liu Y. Recurrent neural networks for multivariate time series with missing values. Sci Rep. 2018;8(1):6085. https://doi.org/10.1038/s41598-018-24271-9.

Pramanik R, Sikdar R, Sarkar R. Transformer-based deep reverse attention network for multi-sensory human activity recognition. Eng Appl Artif Intell. 2023;122:106150. https://doi.org/10.1016/j.engappai.2023.106150.

An Y, Liu Y, Chen X, Sheng Y. TERTIAN: Clinical endpoint prediction in ICU via time-aware transformer-based hierarchical attention network. Comput Intell Neurosci. 2022;2022:4207940.

Tipirneni S, Reddy CK. Self-supervised transformer for sparse and irregularly sampled multivariate clinical time-series. ACM Trans Knowl Discov Data. 2022;16(6):1-27. https://doi.org/10.1145/3531290.

Lim B, Arık SÖ, Loeff N, Pfister T. Temporal fusion transformers for interpretable multi-horizon time series forecasting. Int J Forecast. 2021;37(4):1748-64. https://doi.org/10.1016/j.ijforecast.2021.03.012.

Li J, Wang C, Su W, Ye D, Wang Z. Uncertainty-aware self-attention model for time series prediction with missing values. Fractal Fract. 2025;9(3):181. https://doi.org/10.3390/fractalfract9030181.

van der Steen WE, Leemans EL, van den Berg R, Roos YB, Marquering HA, Verbaan D, et al. Radiological scales predicting delayed cerebral ischemia in subarachnoid hemorrhage: systematic review and meta-analysis. Neuroradiology. 2019;61(3):247-56. https://doi.org/10.1007/s00234-018-2166-1.

Santana LS, Diniz JB, Rabelo NN, Teixeira MJ, Figueiredo EG, Telles JP, et al. Machine learning algorithms to predict delayed cerebral ischemia after subarachnoid hemorrhage: a systematic review and meta-analysis. Neurocrit Care. 2024;40(3):1171-81. https://doi.org/10.1007/s12028-024-01935-2.

Zhang H, Zou P, Luo P, Jiang X. Machine learning for the early prediction of delayed cerebral ischemia in patients with subarachnoid hemorrhage: systematic review and meta-analysis. J Med Internet Res. 2025;27:e54121. https://doi.org/10.2196/54121.

Mohammadzadeh I, Niroomand B, Eini P, Khaledian H, Choubineh T, Luzzi S. Leveraging machine learning algorithms to forecast delayed cerebral ischemia following subarachnoid hemorrhage: a systematic review and meta-analysis of 5,115 participants. Neurosurg Rev. 2025;48(1):26. https://doi.org/10.1007/s10143-024-03311-1.

Shukla SN, Marlin BM. Multi-time attention networks for irregularly sampled time series. arXiv [Preprint]. 2021:arXiv:2101.10318.

Zhao D, Shi Y, Cheng L, Li H, Zhang L, Guo H. Time interval uncertainty-aware and text-enhanced based disease prediction. J Biomed Inform. 2023;139:104239. https://doi.org/10.1016/j.jbi.2023.104239.

Wickstrøm K, Mikalsen KØ, Kampffmeyer M, Revhaug A, Jenssen R. Uncertainty-aware deep ensembles for reliable and explainable predictions of clinical time series. IEEE J Biomed Health Inform. 2021;25(7):2435-44. https://doi.org/10.1109/JBHI.2020.3046516.

Du W, Côté D, Liu Y. SAITS: Self-attention-based imputation for time series. Expert Syst Appl. 2023;219:119619. https://doi.org/10.1016/j.eswa.2023.119619.

Saleh H, El-Sappagh S, McCann M, Alsamhi SH, Breslin JG. Multivariate multi-horizon time-series forecasting for real-time patient monitoring based on cascaded fine tuning of attention-based models. Comput Biol Med. 2025;194:110406. https://doi.org/10.1016/j.compbiomed.2025.110406.

Harerimana G, Kim JW, Jang B. A multi-headed transformer approach for predicting the patient's clinical time-series variables from charted vital signs. IEEE Access. 2022;10:105993-6004. https://doi.org/10.1109/ACCESS.2022.3211662.

Xu Y, Xu S, Ramprassad M, Tumanov A, Zhang C. TransEHR: Self-supervised transformer for clinical time series data. In: Proceedings of Machine Learning for Health (ML4H). PMLR; 2023. p. 623-35.

He R, Chiang JN. Simultaneous forecasting of vital sign trajectories in the ICU. Sci Rep. 2025;15(1):14996. https://doi.org/10.1038/s41598-025-99433-4.

Ma F, Chitta R, Zhou J, You Q, Sun T, Gao J. Dipole: Diagnosis prediction in healthcare via attention-based bidirectional recurrent neural networks. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2017 Aug 13; Halifax, Canada. ACM; 2017. p. 1903-11. https://doi.org/10.1145/3097983.3098088.

Azzam AY, Vaishnav D, Essibayi MA, Unda SR, Jabal MS, Liriano G, et al. Prediction of delayed cerebral ischemia followed aneurysmal subarachnoid hemorrhage: a machine-learning based study. J Stroke Cerebrovasc Dis. 2024;33(4):107553. https://doi.org/10.1016/j.jstrokecerebrovasdis.2024.107553.

Ge S, Chen J, Wang W, Zhang LB, Teng Y, Yang C, et al. Predicting who has delayed cerebral ischemia after aneurysmal subarachnoid hemorrhage using machine learning approach: a multicenter, retrospective cohort study. BMC Neurol. 2024;24(1):177. https://doi.org/10.1186/s12883-024-03683-2.

Liu Y, Li C, Wang H. Machine learning-driven risk prediction of delayed cerebral ischemia after aneurysmal subarachnoid hemorrhage using peripheral inflammatory markers. Front Neurol. 2025;16:1713341. https://doi.org/10.3389/fneur.2025.1713341.

Mao M, Zhou R, Chen Y, Wei J, Lin M, Li W. Construction of a nomogram model for predicting delayed cerebral ischemia in aneurysmal subarachnoid hemorrhage patients. Sci Rep. 2025;15(1):17739. https://doi.org/10.1038/s41598-025-02161-1.

Chen S, Jiang H, He P, Tang X, Chen Q. New grading scale based on early factors for predicting delayed cerebral ischemia in patients with aneurysmal subarachnoid hemorrhage: a multicenter retrospective study. Front Neurol. 2024;15:1393733. https://doi.org/10.3389/fneur.2024.1393733.

Schenck H, van Craenenbroeck C, van Kuijk S, Gommer E, Veldeman M, Temel Y, et al. Systematic review and meta-analysis of transcranial doppler biomarkers for the prediction of delayed cerebral ischemia following subarachnoid hemorrhage. J Cereb Blood Flow Metab. 2025;45(6):1031-47. https://doi.org/10.1177/0271678X251320145.

Xu W, Dai Y, Yang Y, Loza M, Zhang W, Cui Y, et al. OmniTFT: Omni target forecasting for vital signs and laboratory result trajectories in multi center ICU data. arXiv [Preprint]. 2025 Nov 23:arXiv:2511.19485.