

ORIGINAL RESEARCH

Open access

Deep Learning Model for Predicting Inpatient Fall Events in Medical-Surgical Units Using Nursing Progress Notes, Medication Burden, Mobility Assessment Scores, Bed-Exit Alarm Logs, and Room-Level Environmental Risk Indicators

Lucas Meyer^{1*}, Anna Schmid², Stefan Braun¹

Abstract

Inpatient falls in medical-surgical units remain frequent, clinically serious, and difficult to prevent using periodic risk assessment alone. Static scales can support bedside awareness but may miss rapidly changing patient conditions. Existing approaches often fail to integrate nursing narratives, medication burden, mobility scores, bed-exit alarm activity, and room-level environmental hazards. These signals are usually documented in separate systems and are not continuously synthesized into fall risk estimates. This article proposes a multimodal deep learning model to predict the probability of an inpatient fall within the next 24 hours. The model is designed for medical-surgical units and uses both structured and unstructured clinical inputs. The proposed architecture uses a late-fusion design with a clinical text encoder for nursing progress notes and a structured-feature subnetwork for medication burden, mobility assessment scores, alarm logs, and environmental indicators. A final risk-scoring layer would generate a dynamic probability estimate suitable for clinical decision support. Conceptually, the model would produce an updated fall risk score that reflects subtle language cues, recent medication changes, impaired mobility, repeated bed-exit activity, and modifiable room hazards. The score would support continuous surveillance rather than replacing nursing judgment. A multimodal deep learning model could help shift inpatient fall prevention from episodic screening toward continuous, data-driven monitoring. Silent validation and careful workflow integration would be essential before clinical activation.

Keywords Deep learning, Inpatient falls, Clinical informatics, Nursing notes, Medication burden, Bed-exit alarms

*Correspondence:

Lucas Meyer

lucas.meyer@gmail.com

¹ Department of Health Data Science and Digital Systems, Faculty of Engineering, ETH Zurich, Zurich, Switzerland

² Department of Clinical Informatics and Analytics, Faculty of Medicine, University of Bern, Bern, Switzerland

Introduction

Inpatient falls remain a major patient safety concern in acute medical-surgical units because they may result in injury, fear of mobility, prolonged hospitalization, and additional care costs. Prediction models developed from electronic health record data suggest that fall risk is not a static attribute but changes across the hospitalization as clinical status, medications, mobility, and care processes evolve [1, 2]. Clinical informatics studies have therefore increasingly framed fall prediction as a time-sensitive adverse-event prediction task rather than a one-time screening exercise [3, 4]. A model-

oriented approach is especially relevant for medical-surgical settings, where heterogeneous diagnoses and variable nursing workflows create complex risk trajectories [5, 6].

Traditional fall risk scales such as the Morse Fall Scale, Hendrich II, and STRATIFY provide structured bedside assessments but may offer limited adaptability when patient behavior changes within or between shifts. Studies comparing structured risk documentation with machine-learning approaches suggest that relying only on scale scores may miss temporal patterns available in flowsheets, orders, medication records, and nursing documentation [1, 3, 5]. Static tools may also be affected by inter-rater variability and local documentation practices, limiting their portability across units [6, 7]. These limitations motivate a prediction model that treats structured assessment scores as one input stream rather than the sole basis for risk classification.

Hospitals already generate multiple real-time or near-real-time data streams relevant to fall risk, including nursing progress notes, electronic medication administration records, mobility assessments, bed-exit alarm logs, and environmental observations. Nursing notes may describe confusion, impulsivity, toileting attempts, refusal of assistance, or unsteady gait before these conditions are reflected in structured fields [8, 9]. Medication records can capture sedative, antihypertensive, and psychotropic exposure, while alarm systems and room-level factors may provide direct signals of movement and environmental vulnerability [10-13]. Yet these data are often fragmented across clinical systems, making them difficult for nurses to synthesize continuously during routine care.

A multimodal deep learning model could address this fragmentation by combining text-derived risk cues with structured clinical, medication, alarm, and environmental indicators into a continuously updated fall probability. Prior deep learning and multimodal electronic health record studies show the conceptual value of fusing clinical text with structured data for hospital event prediction, while fall-specific work demonstrates the relevance of time-varying EHR signals [14-17]. In this article, the proposed model is not presented as an experimentally validated system but as a conceptual MDL architecture for inpatient fall prediction. Its purpose is to define how nursing documentation, medication burden, mobility assessments, bed-exit alarm activity, and room-level risk indicators could be integrated into a clinically interpretable prediction pipeline.

Background

Inpatient falls: epidemiology and risk factors

Inpatient falls arise from interactions between intrinsic patient factors, such as frailty, cognition, gait instability, toileting needs, and acute illness, and extrinsic factors, such as medications, alarms, room layout, and staffing workflows. Clinical prediction model reviews emphasize that hospital fall risk is multifactorial and context-dependent, making simple rule-based stratification difficult across units and populations [7]. Machine-learning studies using EHR and administrative data further suggest that fall risk is influenced by longitudinal changes in patient condition rather than only admission-level characteristics [2, 3]. Preventive bundles remain important, but model-driven surveillance could help target these bundles when risk is rising during hospitalization [18].

Nursing progress notes as an underexploited signal

Nursing progress notes contain clinically meaningful descriptions of behavior, mobility, cognition, continence, assistance needs, and environmental interactions that may not be fully captured in structured fields. Text mining studies of nursing notes show that fall-relevant information can be extracted from free-text documentation and used to identify patterns associated with risk [9, 19]. Natural language processing approaches are particularly relevant because phrases describing agitation, unsteady gait, impulsive bed exits, or refusal of help may appear before a structured score is updated [8, 20]. For this reason, clinical narrative should be treated as a core input modality in a fall prediction model rather than as supplementary context.

Medication burden and fall risk

Medication burden is a central and potentially modifiable contributor to fall risk, especially when patients receive psychotropics, sedatives, antihypertensives, or multiple fall-risk-increasing drugs. Systematic reviews of cardiovascular and psychotropic medications indicate that drug class, dosage context, and patient vulnerability should be considered together when estimating risk [12, 13]. Geriatric medication safety guidance also emphasizes the need to translate evidence about fall-risk-increasing drugs into clinical workflows that support monitoring and medication review [21]. A predictive model should therefore represent medication burden dynamically, reflecting new administrations, cumulative sedative load, and changes in exposure over time [2].

Bed-exit alarms and environmental monitoring

Bed-exit alarms, pressure-sensitive mats, and related monitoring systems provide movement-related signals that may precede an observed fall, although they can also contribute to alarm fatigue if used indiscriminately. Studies of bed-exit detection and integrated alarm systems suggest that alarm activity can serve as a proxy for restlessness, attempts to mobilize without assistance, or toileting-related movement patterns [10, 11]. Environmental factors such as lighting, clutter, bed height, bed-rail position, floor surfaces, and distance to nursing station can modify risk even when patient-level clinical variables are similar [22]. A model that incorporates alarm patterns and room-level risk indicators could support more context-aware prevention than patient-only screening tools.

Deep learning for multimodal clinical event prediction

Deep learning methods are well suited to multimodal clinical event prediction because they can encode unstructured notes, structured temporal variables, and heterogeneous EHR inputs within a unified architecture. Large-scale EHR deep learning studies demonstrate that clinical prediction can benefit from models that process longitudinal structured data and narrative context together [14, 15]. Clinical language models such as ClinicalBERT provide a foundation for encoding note semantics, while multimodal fusion frameworks show how text and structured information can be combined for downstream prediction [16, 23]. These developments support the conceptual design of a fall prediction model that combines nursing language, medication burden, mobility status, alarm activity, and environmental context [17].

Model Development Overview

High-level predictive pipeline

The proposed pipeline would continuously extract data from the EHR, electronic medication administration record, nursing flowsheets, bed-exit alarm system, and room-environment database, then transform these inputs into synchronized patient-time representations. Fall-specific machine-learning studies support this time-varying approach because risk estimates can change as new clinical documentation, medication administrations, and mobility observations become available [1-3]. The multimodal network would process text and structured inputs separately before combining them into a rolling fall-probability score for the next 24 hours. This pipeline is intended to support surveillance and prioritization, not autonomous clinical decision-making [24, 25].

Core input modalities

The core input modalities would include nursing note free text, a medication burden score, the latest mobility assessment score, recent bed-exit alarm rate, and a room-level environmental risk index. Nursing text would contribute descriptions of cognition, agitation, gait, toileting attempts, and assistance refusal, while structured medication and mobility features would capture measurable risk factors already embedded in routine care [8, 9, 12]. Bed-exit alarms would provide a movement-sensitive signal, and environmental features would contextualize whether the room setup increases vulnerability during unassisted movement [10, 11]. This multimodal specification reflects evidence that fall risk is distributed across clinical, behavioral, pharmacologic, and environmental domains [7, 22].

Design principles

The model should be continuously updating, interpretable for nurses, computationally efficient for real-time inference, and resilient to missing or irregular data. Prior inpatient fall prediction work highlights the importance of time-varying EHR features, while nursing decision support research emphasizes that analytic tools must fit clinical workflow and support rather than obscure judgment [1, 24, 25]. Missingness should be represented explicitly because absent notes, delayed mobility documentation, or unavailable alarm data may reflect workflow patterns rather than true absence of risk. The design should therefore favor modular inputs, transparent outputs, and conservative alerting logic to reduce burden and preserve trust [18].

Data Sources and Feature Engineering

Structured features: medication burden, mobility scores, alarms, and room factors

Structured features would be extracted from the electronic medication administration record, nursing flowsheets, bed-exit alarm logs, and room-level operational databases. Medication burden could summarize active drug count, sedative exposure, psychotropic use, antihypertensive exposure, and recent medication changes, reflecting evidence that fall-risk-increasing drugs require context-sensitive representation [12, 13, 21]. Mobility features would include the most recent documented scale score and indicators of changing assistance needs, while alarm features would represent recent activation frequency and temporal clustering [5, 6, 10]. Room-level indicators could encode lighting adequacy, bed-rail position, bed height, call-light accessibility, clutter risk, and proximity to the nurses' station as modifiable environmental context [11, 22].

Table 1 maps each fall-risk data modality to its predictive meaning, bedside actionability, and potential source of bias or unreliability.

Table 1. Multimodal Fall-Risk Signal Matrix Linking Input Domains to Predictive Meaning, Clinical Actionability, and Bias Risk

Input modality	Predictive contribution	Example fall-risk signals	Clinical actionability	Main bias or reliability concern
Nursing progress notes	Captures subtle behavioral and functional cues before structured fields are updated	Confusion, impulsivity, agitation, refusal of help, toileting attempts, unsteady gait	Supports targeted rounding, supervision, toileting assistance, and handoff prioritization	Documentation timing, subjective wording, under-documentation, variation by nurse
Medication burden	Represents pharmacologic contributors to impaired balance, sedation, hypotension, or cognition	Sedatives, psychotropics, antihypertensives, polypharmacy, recent medication changes	Supports medication review, deprescribing discussion, timing review, monitoring after administration	Dose-response uncertainty, delayed physiologic effects, incomplete medication-risk classification
Mobility assessment scores	Provides structured representation of functional status and assistance needs	Declining mobility score, new assistive-device need, increased transfer assistance	Supports mobility support, physical therapy referral,	Inter-rater variability, delayed reassessment, local scale differences

			assisted ambulation planning	
Bed-exit alarm logs	Captures movement-related instability and unassisted mobilization attempts	Frequent alarms, clustered alarms, nighttime alarms, repeated bed-exit attempts	Supports rounding frequency adjustment, toileting schedule, alarm review, sitter consideration	Alarm fatigue, false alarms, inconsistent device sensitivity, incomplete response linkage
Room-level environmental indicators	Adds modifiable contextual risk beyond patient-level variables	Poor lighting, clutter, bed height, call-light distance, bed-rail position, distance from nursing station	Supports room modification, relocation, low-bed use, environmental safety checks	Inconsistent documentation, site-specific room layouts, limited real-time updates

Unstructured text: nursing progress notes NLP

Unstructured nursing progress notes would be preprocessed into time-stamped clinical text segments, including shift summaries, safety notes, mobility descriptions, fall-risk comments, and withdrawal or delirium-related assessments when present. Context-aware clinical language models could encode these notes into semantic representations that capture fall-relevant concepts such as confusion, impulsivity, gait instability, restlessness, toileting urgency, and refusal of assistance [15, 23]. Fall-specific NLP studies support the idea that nursing narratives contain predictive information not fully reflected in structured variables [8, 9, 19]. The text branch should therefore be fine-tuned conceptually for fall-risk semantics while preserving sufficient interpretability for phrase-level review by clinicians [20].

Multimodal feature alignment and temporal synchronisation

Multimodal alignment would require synchronizing note timestamps, medication administration times, mobility documentation, alarm activity, and room-level observations within clinically meaningful time windows. Because nursing notes may be written once per shift while alarm logs may be continuous and medication records may update at administration events, the model should preserve temporal ordering without assuming equal data cadence [1, 2]. Feature windows should be constructed so that only information available before the prediction time contributes to the fall probability, thereby avoiding forward-looking bias [3, 7]. This temporal synchronization is essential for a model intended to approximate prospective bedside deployment rather than retrospective explanation.

Deep Learning Architecture

Input representation

The model would represent nursing notes as tokenized clinical text sequences and structured variables as normalized vectors with explicit missingness masks. Text input could be encoded using a clinical transformer architecture derived from publicly available clinical language models, while structured features would include medication burden, mobility scores, alarm summaries, and environmental indicators [15, 23]. Missing modalities should not be discarded because the absence of recent mobility scoring, delayed documentation, or unavailable alarm logs may itself carry workflow-relevant information [25]. This input strategy would allow the model to combine dense semantic note embeddings with compact structured representations of evolving fall risk [14].

Multimodal network design

The proposed architecture would include a clinical-text branch that generates a note embedding, a structured-data branch that processes normalized clinical and operational vectors, and a late-fusion layer that combines both representations

before producing a fall probability. Late fusion is appropriate because nursing language and structured risk indicators differ in scale, timing, and meaning, yet each may contribute complementary information [16, 17]. The structured branch could conceptually use fully connected layers with normalization and masking, while the text branch could use transformer-derived embeddings from nursing documentation [15, 23]. This design is consistent with broader multimodal clinical prediction research while remaining tailored to the specific fall-risk domains documented in acute care [14].

Figure 1 illustrates the proposed multimodal deep learning pipeline for integrating nursing narratives, structured clinical indicators, alarm activity, and environmental context into a dynamic 24-hour inpatient fall risk estimate.




Figure 1. Multimodal Deep Learning Pipeline for Dynamic 24-Hour Inpatient Fall Risk Prediction in Medical-Surgical Units

Output: dynamic fall risk score

The output layer would generate a continuously updated fall risk score between zero and one, representing the predicted probability of an inpatient fall within the next 24 hours. The threshold for alerting should be configurable so that hospitals can balance sensitivity, false alerts, nursing workload, and prevention resources rather than relying on a fixed universal cutoff [18, 24]. Risk updates would occur when new notes, medication administrations, mobility scores, bed-exit alarms, or environmental observations become available, allowing the model to reflect intra-shift changes [1, 2]. The score should be displayed with contributing factors so that nurses can interpret the recommendation and connect it to actionable prevention strategies [10, 25].

Addressing Confounding, Imbalance, and Temporal Dynamics

Temporal ordering and window definitions

The prediction task should be defined as estimating fall probability within the next 24 hours using only information available before the prediction timestamp. Features from the preceding 6 to 12 hours could capture recent nursing observations, medication administrations, mobility changes, alarm activity, and room conditions while reducing the risk of forward-looking bias [1, 2]. Encounter boundaries should be respected so that information from prior admissions or post-event documentation does not contaminate the target window [3, 4]. This temporal structure would make the model more consistent with prospective clinical use in medical-surgical units.

Managing severe class imbalance

Because inpatient falls are rare relative to the number of non-fall patient-hours, the model should be developed with explicit attention to severe class imbalance. Conceptually appropriate strategies include class-weighted loss, focal loss, temporal resampling, and threshold selection based on clinically acceptable alert burden rather than accuracy alone [3, 7]. Evaluation should emphasize precision-recall behavior and calibration because a model that appears strong under global discrimination metrics may still generate too many low-value alerts for bedside nurses [18, 24]. This imbalance-aware framing is essential for avoiding a system that increases documentation burden or alarm fatigue without improving prevention.

Temporal validation and concept drift

Temporal validation should train the model on earlier periods and evaluate it on later periods to approximate how performance would behave after deployment. This design is important because fall-prevention practices, documentation habits, medication protocols, alarm use, and unit staffing may change over time [1, 24]. Studies of electronic health record prediction and clinical decision support show that models can be sensitive to workflow and institutional context, making prospective-style validation more informative than random splits [4, 26]. Ongoing monitoring for concept drift would be needed if the model changed nursing behavior or if fall-prevention bundles evolved after implementation.

Model Interpretability and Clinical Trust

Explaining multimodal predictions

The model should provide interpretable explanations that connect the risk score to clinically recognizable factors rather than presenting a black-box probability alone. For nursing notes, attention-weighted or attribution-guided phrase highlighting could identify language related to unsteady gait, confusion, toileting attempts, refusal of assistance, or repeated attempts to get out of bed [8, 9, 19]. For structured inputs, feature-attribution methods could summarize recent sedative exposure, rising medication burden, declining mobility score, frequent bed-exit alarms, or room-level hazards [10, 12, 13]. Explanations should be concise enough for shift workflow while sufficiently transparent to support nurse trust and accountability [25].

Integration into nursing decision support

The prediction output should be integrated into nursing decision support as a dynamic risk summary with top contributing factors and suggested prevention domains. Prior work on fall prediction tools and nursing decision support indicates that analytic outputs are more useful when embedded into existing dashboards, flowsheets, and care-planning routines rather than isolated in separate applications [24, 25, 27]. The model could help charge nurses prioritize rounding, sitter allocation, toileting assistance, bed placement, medication review, and room modification without replacing bedside assessment [18, 21]. Clinical trust would depend on whether the system reduces cognitive burden and supports timely action instead of adding another alert stream.

Clinical Deployment and Operational Integration

Real-time inference at the bedside

For deployment, the model could operate as a clinical microservice consuming real-time or near-real-time data from EHR, medication, nursing documentation, bed-exit alarm, and environmental systems. Interoperability standards such as HL7 or FHIR would support data exchange, but the operational challenge would be ensuring that timestamps, missing values, and delayed documentation are handled consistently [1, 4]. The risk score could appear in the EHR patient summary, nurse handoff view, or nurse call-system interface, with updates triggered by new notes, medication administrations,

alarm events, or mobility documentation [10, 11]. Such integration should be tested silently before activation to understand alert volume and workflow fit.

Tiered preventive interventions

A tiered intervention framework would translate the continuous risk score into clinically meaningful action levels. Low-risk patients might continue standard precautions, medium-risk patients might receive increased rounding, toileting assistance, and environmental review, and high-risk patients might prompt low-bed placement, medication reassessment, alarm review, or one-to-one observation when clinically appropriate [18, 21]. This tiered approach would help prevent the model from being interpreted as a single binary alarm and would align prediction with available prevention resources [10, 24]. It would also allow hospitals to adjust thresholds and interventions based on unit staffing, patient population, and local fall-prevention policy.

Evaluation Strategy

Predictive performance metrics

Evaluation should include discrimination, calibration, and clinical alert-burden metrics without relying on a single performance statistic. Receiver operating characteristic analysis, precision-recall analysis, sensitivity at acceptable false-alert rates, and calibration assessment would together describe whether the model can identify risk while remaining clinically usable [3, 7]. Because falls are uncommon, precision-recall behavior and positive predictive value would be especially important for understanding whether alerts would be actionable in routine nursing workflow [18, 24]. Silent-mode evaluation should also examine whether risk explanations are understandable and linked to plausible prevention actions [25, 27].

Temporal validation across units and hospitals

Validation should begin with internal temporal testing in one health system and then extend to external validation across different medical-surgical units or hospitals. This is necessary because documentation culture, fall-risk scale use, medication practices, alarm systems, room layouts, and patient populations may vary substantially across institutions [4-6]. Cross-site fall prediction studies and interpretable model development work suggest that transportability cannot be assumed even when the same EHR vendor or scale names are used [1, 26]. External validation would therefore test whether the model has learned generalizable risk signals rather than local documentation artifacts.

Clinical utility and workflow impact assessment

Clinical utility should be assessed through prospective silent-mode deployment before any active alerts are shown to nurses. During this phase, investigators could evaluate alert frequency, simulated intervention eligibility, explanation quality, and the extent to which predicted risk aligns with documented clinical concerns, without reporting premature effectiveness claims [24, 25]. Workflow impact assessment should also consider alarm fatigue, nurse acceptance, medication review feasibility, and whether environmental recommendations are actionable in real time [10, 21, 22]. Broader multimodal prediction research supports this staged approach because technically plausible models still require careful evaluation of safety, usability, and clinical integration before activation [16, 17, 28].

Table 2 presents a staged validation and deployment readiness framework for determining whether the proposed model is safe, interpretable, and operationally suitable for clinical use.

Table 2. Validation and Deployment Readiness Framework for a Multimodal Inpatient Fall Prediction Model

Readiness domain	Required evaluation question	Preferred assessment approach	Risk if omitted	Deployment implication
Temporal validity	Does the model predict future falls using only information available before the prediction time?	Prospective-style temporal train-test split; strict prediction windows; leakage audit	Inflated performance from post-event or future information	Model should not proceed to clinical testing without temporal leakage control
Class imbalance handling	Can the model identify rare fall events without excessive false alerts?	Precision-recall analysis, sensitivity at fixed alert burden, calibration by risk tier	High false-alert burden and nurse disengagement	Thresholds must be selected around workflow tolerance, not accuracy alone
Multimodal contribution	Do text, medication, mobility, alarm, and room features add complementary value?	Ablation analysis comparing single-modality and fused models	Complex model may add burden without meaningful gain	Retain only modalities that improve prediction, explanation, or actionability
Interpretability	Can nurses understand why the patient is flagged as higher risk?	Phrase-level note attribution, structured feature contribution summaries, clinician review	Black-box risk scores may reduce trust and accountability	Explanations must be concise, clinically recognizable, and linked to prevention
Transportability	Does performance hold across units, hospitals, documentation styles, and room layouts?	External validation across medical-surgical units and health systems	Model may learn local documentation artifacts rather than generalizable risk	Local calibration and multi-site testing are required before broad use
Workflow safety	Does the system support prevention without worsening alarm fatigue or documentation burden?	Silent-mode deployment, simulated alert review, usability testing, nurse feedback	Increased cognitive load, ignored alerts, or inappropriate intervention escalation	Active alerts should begin only after silent validation and workflow review
Governance and monitoring	Does performance remain stable after practice patterns or fall-prevention policies change?	Drift monitoring, recalibration plan, periodic safety review, audit logs	Degraded performance over time and unsafe reliance on outdated risk estimates	Ongoing monitoring is required after deployment

Limitations

Data quality and documentation bias

The proposed model would depend on the quality, timing, and completeness of nursing documentation, medication records, mobility scoring, alarm logs, and environmental observations. Nursing notes may omit relevant bedside observations, structured scores may be delayed, and medication administration timestamps may not perfectly represent physiologic effect onset. Bed-exit alarms may vary in sensitivity, staff response, and documentation linkage, while room-

level risk indicators may be inconsistently recorded. These limitations mean that the model should be viewed as a decision-support tool requiring prospective validation rather than as a replacement for bedside judgment.

Generalizability across hospitals and patient populations

A model developed in one health system may not transfer directly to another hospital with different documentation norms, fall-prevention workflows, medication formularies, mobility scales, alarm technologies, or patient demographics. Medical-surgical units also vary in case mix, staffing ratios, room design, and availability of sitters or mobility support. These differences could alter both the meaning of model inputs and the feasibility of recommended interventions. Multi-site validation and local calibration would therefore be necessary before broad deployment.

Conclusion

The proposed MDL framework describes a multimodal deep learning model for predicting inpatient fall events in medical-surgical units. It integrates nursing progress notes, medication burden, mobility assessment scores, bed-exit alarm logs, and room-level environmental indicators into a continuously updated fall risk estimate. The model is intended to support clinical surveillance and prevention planning rather than replace nursing assessment. Its central contribution is a conceptual architecture for combining fragmented fall-risk signals into a unified predictive workflow.

A key strength of this approach is its ability to combine unstructured nursing narratives with structured clinical and operational data. Nursing notes can capture subtle behavioral and mobility cues, while medication records, mobility scores, alarms, and environmental indicators provide complementary risk information. Continuous updating would allow the model to respond to intra-shift changes that static assessment scales may miss. This makes the framework especially relevant for dynamic medical-surgical environments.

Important challenges remain before such a model could be safely used in practice. Data quality, documentation bias, severe class imbalance, temporal drift, and institutional variation could all affect reliability. Interpretability and workflow integration would be as important as model design because nurses must understand and trust the risk estimate. Prospective validation would be required to determine whether the system supports timely, feasible, and equitable fall-prevention actions.

Future work should begin with silent pilot deployments that estimate alert burden, explanation usefulness, and operational fit without changing care. Multi-site studies should then evaluate whether the model generalizes across hospitals and whether risk-guided interventions can reduce preventable falls without worsening alarm fatigue. The ultimate goal should be a carefully governed decision-support system that improves patient safety while respecting nursing expertise. Broad clinical activation should occur only after rigorous validation, usability testing, and assessment of patient-centered impact.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 05 Sep 2021

Revised: 29 Oct 2021

Accepted: 22 Nov 2021

Published online: 25 February 2022

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

Cho I, Boo EH, Chung E, Bates DW, Dykes P. Novel approach to inpatient fall risk prediction and its cross-site validation using time-variant data. *J Med Internet Res.* 2019;21(2):e11505.

<https://doi.org/10.2196/11505>.

Choi Y, Staley B, Henriksen C, Xu D, Lipori G, Brumback B, et al. A dynamic risk model for inpatient falls. *Am J Health Syst Pharm.* 2018;75(17):1293-303.

<https://doi.org/10.2146/ajhp180013>.

Lindberg DS, Prosperi M, Bjarnadottir RI, Thomas J, Crane M, Chen Z, et al. Identification of important factors in an inpatient fall risk prediction model to improve the quality of care using EHR and electronic administrative data: A machine-learning approach. *Int J Med Inform.* 2020;143:104272.

<https://doi.org/10.1016/j.ijmedinf.2020.104272>.

Moskowitz G, Egorova NN, Hazan A, Freeman R, Reich DL, Leipzig RM. Using electronic health records to enhance predictions of fall risk in inpatient settings. *Jt Comm J Qual Patient Saf.* 2020;46(4):199-206.

<https://doi.org/10.1016/j.jcjq.2019.12.003>.

Yokota S, Endo M, Ohe K. Establishing a classification system for high fall-risk among inpatients using support vector machines. *Comput Inform Nurs.* 2017;35(8):408-16.

<https://doi.org/10.1097/CIN.0000000000000346>.

Jung H, Park HA, Hwang H. Improving prediction of fall risk using electronic health record data with various types and sources at multiple times. *Comput Inform Nurs.* 2020;38(3):157-64.

<https://doi.org/10.1097/CIN.0000000000000582>.

Parsons R, Cramb SM, McPhail SM. Clinical prediction models for hospital falls: a scoping review protocol. *BMJ Open.* 2021;11(9):e051047.

<https://doi.org/10.1136/bmjopen-2021-051047>.

Nakatani H, Nakao M, Uchiyama H, Toyoshiba H, Ochiai C. Predicting inpatient falls using natural language processing of nursing records obtained from Japanese electronic medical records: case-control study. *JMIR Med Inform.* 2020;8(4):e16970.

<https://doi.org/10.2196/16970>.

Bjarnadottir RI, Lucero RJ. What can we learn about fall risk factors from EHR nursing notes? A text mining study. *eGEMS*. 2018;6(1):21.
<https://doi.org/10.5334/egems.237>.

Seow JP, Chua TL, Aloweni F, Lim SH, Ang SY. Effectiveness of an integrated three-mode bed exit alarm system in reducing inpatient falls within an acute care setting. *Jpn J Nurs Sci*. 2022;19(1):e12446.
<https://doi.org/10.1111/jjns.12446>.

Jähne-Raden N, Kulau U, Marschollek M, Wolf KH. INBED: a highly specialized system for bed-exit-detection and fall prevention on a geriatric ward. *Sensors (Basel)*. 2019;19(5):1017.
<https://doi.org/10.3390/s19051017>.

de Vries M, Seppala LJ, Daams JG, van de Glind EMM, Masud T, van der Velde N; EUGMS Task and Finish Group on Fall-Risk-Increasing Drugs. Fall-Risk-Increasing Drugs: A Systematic Review and Meta-Analysis: I. Cardiovascular Drugs. *J Am Med Dir Assoc*. 2018;19(4):371.e1-371.e9.
<https://doi.org/10.1016/j.jamda.2017.12.013>.

Hart LA, Phelan EA, Yi JY, Marcum ZA, Gray SL. Use of fall risk-increasing drugs around a fall-related injury in older adults: a systematic review. *J Am Geriatr Soc*. 2020;68(6):1334-43.
<https://doi.org/10.1111/jgs.16369>.

Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Hardt M, et al. Scalable and accurate deep learning with electronic health records. *NPJ Digit Med*. 2018;1(1):18.
<https://doi.org/10.1038/s41746-018-0029-1>.

Huang K, Altosaar J, Ranganath R. ClinicalBERT: modeling clinical notes and predicting hospital readmission. *arXiv [Preprint]*. 2019;arXiv:1904.05342.

Huang SC, Pareek A, Seyyedi S, Banerjee I, Lungren MP. Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines. *NPJ Digit Med*. 2020;3(1):136.
<https://doi.org/10.1038/s41746-020-00341-z>.

Li Y, Mamouei M, Salimi-Khorshidi G, Solares JRA, Rao R, Hassey A, et al. Hi-BEHT: hierarchical transformer-based model for accurate prediction of clinical events using multimodal longitudinal electronic health records. *IEEE J Biomed Health Inform*. 2023;27(2):1106-17.

Jellett J, Williams C, Clayton D, Plummer V, Haines T. Falls risk score removal does not impact inpatient falls: a stepped-wedge, cluster-randomised trial. *J Clin Nurs*. 2020;29(23-24):4505-13.
<https://doi.org/10.1111/jocn.15467>.

Topaz M, Murga L, Gaddis KM, McDonald MV, Bar-Bachar O, Goldberg Y, ET AL. Mining fall-related information in clinical notes: Comparison of rule-based and novel word embedding-based machine learning approaches. *J Biomed Inform*. 2019;90:103103.
<https://doi.org/10.1016/j.jbi.2019.103103>.

Kawazoe Y, Shimamoto K, Shibata D, Shinohara E, Kawaguchi H, Yamamoto T. Impact of a clinical text-based fall prediction model on preventing extended hospital stays for elderly inpatients: model development and performance evaluation. *JMIR Med Inform*. 2022;10(7):e37913.
<https://doi.org/10.2196/37913>.

Seppala LJ, van der Velde N, Masud T, Hartikainen S, Mol A, Strandberg T, et al. EuGMS task and finish group on fall-risk-increasing drugs (FRIDs): position on knowledge dissemination, management, and future research. *Eur Geriatr Med*. 2019;10(2):275-83.
<https://doi.org/10.1007/s41999-019-00162-8>.

Usmani S, Saboor A, Haris M, Khan MA, Park H. Latest research trends in fall detection and prevention using machine learning: a systematic review. *Sensors (Basel)*. 2021;21(15):5134.
<https://doi.org/10.3390/s21155134>.

Alsentzer E, Murphy J, Boag W, Weng WH, Jindi D, Naumann T, et al. Publicly available clinical BERT embeddings. In: Proceedings of the 2nd Clinical Natural Language Processing Workshop. Minneapolis, MN, USA: Association for Computational Linguistics; 2019. p.72-8.

Cho I, Jin IS, Park H, Dykes PC. Clinical impact of an analytic tool for predicting the fall risk in inpatients: controlled interrupted time series. *JMIR Med Inform*. 2021;9(11):e26456.
<https://doi.org/10.2196/26456>.

Akbar S, Lyell D, Magrabi F. Automation in nursing decision support systems: a systematic review of effects on decision making, care delivery, and patient outcomes. *J Am Med Inform Assoc*. 2021;28(11):2502-13.

Shim S, Yu JY, Jekal S, Kim JH, Lee YS, Park H, et al. Development and validation of interpretable machine learning models for inpatient fall events and electronic medical record integration. *Clin Exp Emerg Med*. 2022;9(4):345-53.

Lytle KS, Westra BL, Whittenburg L, Adams M, Akre M, Ali S, et al. Information Models Offer Value to Standardize Electronic Health Record Flowsheet Data: A Fall Prevention Exemplar. *J Nurs Scholarsh*. 2021;53(3):306-14.
<https://doi.org/10.1111/jnu.12646>.

Golmaei SN, Luo X. DeepNote-GNN: predicting hospital readmission using clinical notes and patient network. In: Proceedings of the 12th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics. New York, NY: ACM; 2021. p.1-9.
<https://doi.org/10.1145/3459930.3469511>.