

ORIGINAL RESEARCH

Open access

Sequence Learning Model for Predicting Repeat Diagnostic Imaging Orders Using Prior Imaging History, Symptom Documentation, Specialist Recommendations, Ordering Physician Behavior, and Recent Test Results

Omar Khalid^{1*}, Sara Nadeem¹, Bilal Farooq², Hina Saeed¹

Abstract

Repeat diagnostic imaging is a major driver of healthcare costs, radiation exposure, workflow burden, and downstream follow-up from incidental findings. Its occurrence often follows recognizable patterns shaped by prior imaging, persistent symptoms, specialist advice, recent results, and provider ordering habits. Current utilization management tools often respond after an imaging order has already been placed or rely on static appropriateness rules. They therefore miss opportunities to anticipate repeat ordering risk before the clinician reaches the final order-entry step. This manuscript proposes a sequence learning model that could predict the probability of a repeat diagnostic imaging order within a future clinical window. The model would integrate prior imaging history, symptom documentation, specialist recommendations, ordering physician behavior, and recent test results. The conceptual model would use a recurrent neural network or Transformer encoder to process temporally ordered imaging and clinical events. Structured radiology information system data would be combined with natural language processing features from clinical notes and consult documentation, physician-level ordering context, and recent laboratory or imaging-result signals. Conceptually, the model would forecast whether a repeat CT, MRI, ultrasound, or related diagnostic imaging order is likely to occur. It would also identify major contextual drivers so that the prediction could support pre-emptive review, alternative care suggestions, or guideline-aligned follow-up. A sequence learning model for repeat imaging prediction could function as a safety-and-value layer within radiology workflow. It could reduce unsupported imaging variation while preserving clinically indicated follow-up and surveillance.

Keywords Clinical decision support, Natural language processing, Electronic health records, Sequence learning, Repeat imaging, Radiology utilization

*Correspondence:

Omar Khalid
omar.khalid@gmail.com

¹ Department of Digital Healthcare Systems, Faculty of Medicine, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia

² Department of Clinical Informatics Engineering, Faculty of Engineering, Qatar University, Doha, Qatar

Introduction

Repeat diagnostic imaging creates a layered burden that extends beyond the immediate cost of a single examination. Additional CT, MRI, ultrasound, or radiographic studies can expose patients to radiation, contrast risk, incidental findings, downstream procedures, and anxiety, while also consuming limited scanner capacity and radiology reporting time.

Reviews of low-value imaging emphasize that unnecessary or weakly indicated studies remain common across health systems, especially when prior examinations, indications, and follow-up recommendations are not visible at the point of care [1-3]. A predictive model that recognizes when a repeat order is likely could therefore address a clinically important utilization problem before it becomes an administrative or safety issue.

Most current approaches to imaging overuse rely on radiology benefit managers, clinical decision support, retrospective audit, or educational feedback, each of which has temporal limitations. Decision support may be triggered only after the order has been selected, whereas audit and feedback often occur weeks or months after the ordering behavior has already taken place. Studies of imaging decision support and low-value imaging interventions suggest that such tools can influence appropriateness, but their effect depends on workflow timing, clinician acceptance, and the relevance of guidance to the specific patient context [4-6]. A sequence-aware prediction layer would shift the task from reacting to an order toward anticipating which patient-clinician encounters are moving toward repeat imaging.

The growing availability of radiology information system data, electronic health record timelines, free-text clinical notes, and provider-level utilization patterns makes repeat imaging a suitable target for deep learning. Natural language processing can extract symptoms, indications, and follow-up advice from clinical documentation and radiology reports, while structured data can represent exam type, body region, ordering clinician, recent laboratory values, and prior imaging findings [7-9]. Sequence models developed for electronic health records demonstrate how temporally ordered events can be converted into patient-state representations that support clinical event prediction [10-12]. In imaging utilization, these methods could connect prior imaging history, current clinical narrative, and physician behavior into a unified prediction of future repeat orders.

The central thesis of this article is that a sequence learning model could forecast repeat diagnostic imaging orders before they occur, enabling proactive appropriateness review without relying only on static rules. Such a model would not decide whether imaging is correct; rather, it would estimate the probability of a repeat order and explain the features that make the order likely. Prior work on protocol selection, radiology report mining, clinician variation, and clinical decision support shows that imaging workflows already contain data elements suitable for model-driven assistance. The proposed model therefore frames repeat imaging prediction as a temporal, contextual, and clinician-facing task rather than a purely administrative utilization screen.

Background

Repeat imaging: definitions, causes, and trends

Repeat imaging can include clinically appropriate surveillance, guideline-driven follow-up, unresolved symptom evaluation, duplicate testing after care transitions, and low-value repetition of recent studies. The distinction between necessary follow-up and unnecessary duplication is difficult because the same event, such as a repeat CT, may be appropriate for cancer surveillance but questionable for a stable symptom with a recent negative study. Reviews of low-value imaging show that unnecessary imaging is not caused by a single mechanism; it reflects patient expectations, diagnostic uncertainty, fragmented records, specialty norms, and system incentives [1, 2, 13]. A sequence learning approach should therefore represent repeat imaging not as inherently inappropriate but as an event whose meaning depends on prior results, timing, symptoms, recommendations, and clinical context.

Prior imaging history as a predictor of future imaging

Prior imaging history is one of the most direct predictors of future imaging because it captures both the patient's diagnostic trajectory and the clinician's awareness of unresolved or monitored findings. A patient who has undergone repeated chest CT for pulmonary findings, serial MRI for neurologic symptoms, or ultrasound follow-up after an indeterminate abnormality has a structured timeline that can signal whether another order is expected. Radiology-focused machine learning applications have shown that imaging metadata, report content, indications, and ordering context can

support protocol selection, follow-up identification, and workflow prioritization [8, 14, 15]. In the proposed model, prior imaging events would function as the backbone sequence from which temporal recurrence, modality switching, and body-region persistence are learned.

Symptom documentation and specialist recommendations

Free-text symptom documentation and specialist recommendations provide the clinical rationale behind repeat imaging orders, especially when structured diagnosis codes are too broad or delayed. Natural language processing methods in radiology have been used to identify actionable reports, follow-up recommendations, critical findings, and protocol-relevant clinical indications, demonstrating the feasibility of extracting order-driving cues from narrative text [7, 9, 16, 17]. Consultant notes may also contain explicit recommendations for repeat imaging, while primary care or emergency notes may show persistent symptoms that make repeat imaging more likely. A model that ignores these narratives would risk treating every repeat order as a pattern of utilization rather than as a response to evolving clinical concern.

Ordering physician behavior and practice variation

Ordering physician behavior is a major contextual factor because clinicians differ in their thresholds for diagnostic certainty, their familiarity with appropriateness criteria, their response to patient expectations, and their specialty-specific norms. Emergency department and ambulatory imaging studies have shown that individual physicians and practice settings can contribute to variation in imaging use, even when patient populations are similar [5, 18]. Low-value imaging research further suggests that clinician culture, habit, perceived medicolegal risk, and local workflow constraints shape whether imaging is ordered or deferred [13, 19]. Incorporating physician-level features would allow the proposed model to distinguish patient-driven repeat imaging risk from provider-specific ordering tendencies.

Sequence learning for clinical event forecasting

Sequence learning models are well suited to clinical event prediction because patient histories unfold as irregular, multimodal, and temporally ordered records. Recurrent neural networks, Transformer encoders, and related architectures have been used to represent longitudinal electronic health records, model patient timelines, and predict future clinical events from structured and narrative data [10–12, 20]. More recent generative and Transformer-based approaches show that patient trajectories can be encoded as sequences of diagnoses, procedures, medications, results, and encounters, enabling models to learn long-range dependencies rather than isolated risk factors [21, 22]. For repeat imaging, the same principle can be applied to a timeline in which imaging events, symptoms, specialist advice, and test results jointly shape the probability of a future order.

Model Development Overview

High-level predictive pipeline

The proposed pipeline would begin by extracting each patient's imaging timeline from the radiology information system and electronic health record, including modality, body region, indication, ordering clinician, timestamp, and available report-derived findings. These events would be fused with natural language processing outputs from symptom documentation and specialist recommendations, recent laboratory or imaging-result signals, and physician identifiers before being passed to a sequence model. Prior work on radiology protocol automation, follow-up recommendation detection, and clinical timeline modelling supports the feasibility of converting heterogeneous imaging and EHR data into prediction-ready representations [8, 11, 15]. The model would output a probability of repeat imaging within a clinically relevant future interval, with interpretation features designed for use during order entry or pre-order review.

Core input characteristics

The core input would be a longitudinal sequence of prior imaging events, enriched by exam type, body part, order indication, report impression, follow-up language, and elapsed time since previous studies. Symptom descriptions from clinical notes and specialist consult text would be embedded as contextual signals, while static or slowly changing features such as physician specialty, historical ordering tendency, and recent abnormal test results would be appended to the patient representation. Natural language processing studies in radiology demonstrate that clinical indications and follow-up recommendations can be extracted from narrative documentation, while physician-variation studies show that ordering behavior is a meaningful contextual signal [7, 8, 17, 18]. Together, these inputs would let the model represent repeat imaging as the product of patient trajectory, textual clinical rationale, and provider practice pattern.

Design principles

The model should be sequential, context-aware, dynamically updated, interpretable, and deployable within real-time clinical workflow. Sequential design would preserve the order and spacing of prior events, while contextual design would prevent the model from treating appropriate surveillance and unsupported duplication as equivalent. Interpretability is especially important because clinical decision support tools are more likely to be accepted when they provide understandable reasons rather than unexplained alerts [4-6]. Deployment should therefore emphasize soft guidance, transparent drivers, and integration with radiology-specific appropriateness review rather than replacing clinician judgment.

Figure 1 presents the proposed end-to-end sequence learning workflow for predicting repeat diagnostic imaging orders before order completion and translating model outputs into human-reviewed radiology utilization support.




Figure 1. End-to-end sequence learning workflow for predicting repeat diagnostic imaging orders before order completion

Data Sources and Feature Engineering

Constructing the imaging history sequence

Constructing the imaging history sequence would require extracting exam codes, modality, body region, timestamps, indications, report availability, and order context from radiology information systems and electronic health record order tables. Each imaging event would become a time step, allowing the model to learn whether repeated studies cluster around unresolved findings, follow-up recommendations, chronic disease surveillance, emergency visits, or provider habit. Machine learning applications in radiology have already used structured imaging metadata and workflow information for protocol selection and operational decision support, indicating that such features can be engineered into usable prediction inputs [14, 15, 23, 24]. The sequence should also encode absence of recent imaging when relevant, because a repeat after a long interval has a different meaning from a duplicate study shortly after a prior negative examination.

Extracting symptom and specialist recommendation features

Symptom and recommendation features would be extracted from encounter notes, emergency department documentation, specialist consults, and radiology reports using natural language processing. These features could identify persistent symptoms such as cough, headache, abdominal pain, or back pain, as well as explicit recommendations such as repeat imaging after an interval or additional imaging for an indeterminate finding. Prior work on radiology natural language processing, actionable report identification, and follow-up recommendation detection shows how free-text documentation can be converted into clinically meaningful signals for downstream workflow support [8, 9, 16, 17]. In the proposed model, such features would help distinguish repeat imaging driven by documented clinical progression from repetition driven mainly by uncertainty, habit, or fragmented information.

Physician behavior and recent result encoding

Physician behavior features would encode specialty, care setting, historical ordering pattern, prior response to decision support, and relative tendency to order repeat imaging within similar clinical contexts. Recent result features would represent abnormal laboratory values, new imaging findings, unresolved impressions, or documented follow-up advice that might legitimately increase the probability of another imaging order. Evidence on physician-level imaging variation and decision support response indicates that provider behavior can materially shape utilization, while radiology follow-up studies show that report-derived findings and recommendations influence subsequent imaging pathways [4, 8, 18, 25]. Encoding these variables would help the model separate patient-specific clinical need from clinician-specific ordering propensity.

Table 1 summarizes the major input domains, temporal representations, predictive functions, and operational interpretations used in the proposed repeat imaging sequence model.

Table 1. Input domains, temporal representations, and predictive role of variables in the proposed repeat imaging sequence model

Input domain	Example variables or data elements	Temporal representation in the model	Primary predictive role	Operational interpretation
Prior imaging history	Modality, body region, exam timestamp, indication, ordering location, prior report availability	Sequential imaging-event tokens ordered by time with elapsed-time encoding	Establishes whether the patient has a recent, recurrent, or surveillance-based imaging trajectory	Helps distinguish rapid duplication from expected follow-up or long-interval reassessment
Imaging report content	Impression text, unresolved abnormality, indeterminate finding, recommendation language, comparison with prior study	NLP-derived features attached to imaging-event tokens	Identifies whether prior imaging created a clinically plausible reason for another study	Prevents the model from treating all repeat imaging as inappropriate
Symptom documentation	Persistent headache, abdominal pain, back pain, cough, neurologic symptoms, worsening	Embedded note-derived features linked to encounter timing	Captures the clinical rationale that may precede repeat ordering	Helps separate evolving clinical concern from unsupported repetition

	symptoms, unchanged symptoms			
Specialist recommendations	Consult note recommendation, surveillance interval, suggested modality, recommended follow-up timing	Recommendation tokens or contextual features aligned to patient timeline	Signals planned or specialist-directed imaging pathways	Reduces inappropriate flagging of guideline-concordant follow-up
Recent test results	Abnormal laboratory values, new imaging finding, unresolved radiology impression, recent negative study	Structured result features appended to recent time steps or patient-context vector	Indicates whether new clinical evidence supports or weakens the need for repeat imaging	Supports contextual review rather than volume reduction alone
Ordering physician behavior	Specialty, care setting, historical repeat-order rate, prior response to alerts, ordering tendency within similar cases	Static or slowly changing provider-context vector	Captures clinician-level practice variation and ordering thresholds	Helps identify variation patterns without equating clinician behavior with inappropriate care
Encounter and care setting context	Emergency, inpatient, ambulatory, post-discharge, specialty clinic, transfer of care	Encounter-context feature attached to each event or prediction episode	Accounts for workflow-specific risk of duplicate or repeat imaging	Supports site-specific and setting-specific deployment rules
Patient clinical context	Comorbidities, cancer surveillance status, chronic disease burden, prior abnormal findings, high-risk history	Patient-context vector combined with sequence representation	Personalizes repeat-order risk to underlying clinical need	Protects high-risk patients from inappropriate suppression of necessary follow-up
Future prediction window	Repeat order within defined near-term or medium-term interval	Binary or time-windowed prediction target	Defines whether the model predicts imminent duplicate ordering or later surveillance-related repeat imaging	Allows the system to align prediction timing with CPOE review or pre-order planning
Human-facing explanation layer	Most influential prior imaging event, recommendation phrase, symptom pattern, abnormal result, ordering-context signal	Post-model attribution or temporal contribution summary	Makes the prediction reviewable by clinicians and radiology governance teams	Supports acceptance, modification, or dismissal of the model-generated nudge

Sequence Learning Architecture

Sequence input representation

Each time step in the input sequence would represent either an imaging event, a clinical documentation event, or an extracted recommendation relevant to future imaging. The event embedding would combine modality, body region, indication, report-derived finding, elapsed time, and note-derived symptom or recommendation features, while static context such as ordering physician profile and recent laboratory status would be appended or injected through a context vector. Deep EHR and BEHRT-style representations demonstrate how heterogeneous clinical events can be embedded into temporally ordered patient sequences for prediction tasks [10-12]. For repeat imaging prediction, this representation would allow the model to learn not only that imaging occurred, but also why it occurred, who ordered it, and how recently related evidence appeared.

Recurrent or transformer encoder

The encoder could be implemented as a recurrent neural network, gated recurrent unit, temporal convolutional model, or Transformer-based architecture, depending on deployment needs and the desired handling of long-range dependencies. Recurrent models would be conceptually attractive for sequential patient histories, whereas Transformer encoders could attend across distant imaging events, specialist notes, and result changes that jointly influence repeat imaging probability. Transformer-based EHR models and patient-timeline systems show that attention mechanisms can represent longitudinal clinical context across many event types, making them suitable for modelling imaging trajectories that unfold over repeated encounters [20-22]. The encoder would therefore learn patterns that distinguish expected follow-up from excessive or poorly supported repetition.

Output layer and decision threshold

The output layer would produce a repeat-imaging probability for a future clinical window, using a sigmoid activation or comparable binary risk formulation. The threshold for triggering guidance should be chosen with clinical governance input, emphasizing sensitivity to potentially unnecessary repeats while preserving clearly indicated follow-up for surveillance, worsening symptoms, or specialist-directed care. Low-value imaging intervention studies indicate that overuse reduction must be balanced against appropriateness, clinician trust, and patient safety rather than pursued as a simple volume-minimization goal [1, 3, 19]. For this reason, the output should be interpreted as a prompt for contextual review rather than a determination that the order is inappropriate.

Modelling Temporal Dependencies and Contextual Factors

Time-aware sequence construction

Time-aware sequence construction would encode not only the order of prior events but also the interval between them. A scan performed yesterday, a scan performed three months ago, and a scan performed several years ago should contribute differently to repeat-imaging risk because clinical urgency, duplication risk, and follow-up appropriateness depend on recency. Emergency department CT prediction work illustrates how clinical state at the time of care can be modelled from available EHR signals, while longitudinal EHR models show that temporal spacing is central to patient-event forecasting [20, 26]. In this framework, inter-event duration would help the model distinguish rapid duplication from planned surveillance or delayed reassessment.

Personalizing with patient-level context

Patient-level context would be needed to prevent the model from over-flagging clinically appropriate repeat imaging. Chronic illness, cancer surveillance, recurrent symptoms, prior abnormal findings, and comorbidity patterns may make repeat studies expected rather than excessive. Representation-learning studies using multimodal electronic health records suggest that patient phenotypes can be modelled from longitudinal histories and comorbidities, allowing predictions to reflect the broader clinical state rather than isolated events [12, 27, 28]. For repeat imaging, personalization would help the model treat a surveillance MRI in a high-risk patient differently from a duplicate study in a low-risk patient with a recent negative examination.

Regularization against over-wary flagging

The model should be regularized against over-wary flagging so that it does not treat all repeat imaging as low value. Training labels and governance rules would need to account for appropriate repeats, such as guideline-concordant follow-up, specialist-directed reassessment, and imaging after meaningful clinical change. Low-value imaging research emphasizes that the goal is not indiscriminate reduction, but better distinction between warranted and unwarranted use across clinical contexts [1, 2, 29]. Conceptually, the model should therefore learn patterns of unsupported variability while preserving the pathway for clinically justified imaging.

Interpretability for Appropriate Use Feedback

Explaining repeat predictions during order entry

Interpretability would be essential because clinicians are unlikely to trust a repeat-imaging alert without understanding the basis for the prediction. Attention weights, temporal contribution summaries, or feature-attribution methods could highlight that the model's prediction is driven by a recent similar CT, a specialist recommendation, persistent symptom documentation, or an abnormal prior finding. Radiology natural language processing and follow-up recommendation studies show that report and note content can be transformed into clinically meaningful explanatory signals rather than opaque text embeddings alone [8, 9, 17]. In practice, the explanation should connect the predicted repeat order to recognizable clinical events so that the ordering physician can accept, modify, or dismiss the guidance.

Integrated decision support and clinician nudges

Integrated decision support should surface the prediction as a soft nudge rather than a hard stop. A risk score could be accompanied by concise context, such as the most recent comparable imaging study, relevant follow-up language, and a guideline link when available. Prior imaging decision support and utilization-management studies suggest that clinician response depends on timing, clarity, and perceived relevance, not simply on the presence of an alert [4-6]. The model should therefore support shared review at the order-entry moment while avoiding excessive interruption of urgent or clearly indicated care.

Clinical Integration and Decision Support

Embedding in CPOE with radiology-specific alerts

The model could be embedded in computerized provider order entry so that it is invoked when a clinician selects an imaging order. It would query the patient's recent imaging timeline, symptom documentation, specialist notes, recent results, and provider context, then return a repeat-imaging probability with a short explanation. Radiology protocol automation and clinical decision support studies demonstrate that imaging-order workflows can accommodate algorithmic assistance when the output is operationally relevant and aligned with radiology practice [15, 23, 24]. Such integration would position the model as a real-time appropriateness aid rather than a retrospective utilization report.

Post-order audit and feedback to departments

Beyond real-time order entry, aggregated model outputs could support post-order audit and feedback for departments, service lines, or clinician groups. The system could identify patterns such as frequent repeat imaging after recent negative studies, variation by specialty, or recurring lack of documented indication before repeat orders. Prior work on physician-level variation and low-value imaging interventions suggests that audit and feedback may be most useful when it is targeted, contextual, and connected to local practice norms [13, 18, 19]. Departmental review could therefore use model

explanations to guide education, workflow redesign, and appropriateness discussions rather than merely ranking clinicians by volume.

Evaluation Strategy

Predictive performance metrics

Evaluation should examine discrimination, calibration, and clinical usefulness without treating performance as the only success criterion. Conceptually, AUROC could assess repeat-order classification, precision-recall analysis could focus on the repeat-imaging class, and calibration could determine whether predicted probabilities align with observed ordering risk across patient groups. Machine learning studies in imaging utilization and EHR prediction show that model evaluation should consider workflow-specific targets, temporal drift, and the consequences of false alerts [21, 25, 26]. The model should also be assessed across multiple future windows because a near-term duplicate order and a later surveillance study represent different clinical scenarios.

Temporal and external validation

Temporal validation would test the model on future patient encounters after training, helping determine whether it remains useful as ordering culture, documentation habits, and imaging guidelines change. External validation at another institution would be important because radiology availability, specialist referral patterns, decision support infrastructure, and local norms can differ substantially across health systems. Model-drift work in diagnostic imaging follow-up and broader EHR modelling studies underscore that temporal stability cannot be assumed when clinical workflows evolve [20, 22, 25]. A credible evaluation strategy should therefore include forward-time testing and cross-site assessment before operational deployment.

Prospective effect on imaging volume and appropriateness

A prospective evaluation could examine whether model-supported decision support changes repeat imaging volume, appropriateness review patterns, and clinician behavior. A randomized, pragmatic, or stepped-wedge design would be conceptually appropriate because the intervention affects workflow, departmental norms, and ordering decisions rather than only individual predictions. Studies of low-value imaging reduction, decision support, and low back pain imaging interventions indicate that implementation impact should be judged by appropriateness, clinician acceptance, and unintended consequences as well as utilization change [1, 5, 19, 29]. The evaluation should therefore ask whether the model reduces unsupported repeat imaging while maintaining timely follow-up for patients who need it.

Table 2 outlines the evaluation, interpretability, governance, and implementation safeguards needed before the proposed repeat imaging prediction model could be safely deployed in radiology workflow.

Table 2. Evaluation, interpretability, governance, and implementation framework for safe deployment of the repeat imaging prediction model

Evaluation or governance domain	Recommended assessment approach	Why it matters for this manuscript	Practical deployment implication
Discrimination	AUROC, precision-recall analysis, sensitivity, specificity, positive predictive value, negative predictive value	Determines whether the model can identify encounters likely to generate repeat imaging orders	Should be reported separately for near-term duplicate risk and longer-window follow-up risk

Calibration	Calibration plots, calibration intercept and slope, Brier score, subgroup calibration	The model output is intended to be a probability, not only a rank score	Poor calibration could cause excessive alerts or under-identification of high-risk repeat ordering
Temporal validation	Train on earlier encounters and test on later encounters	Ordering culture, documentation practice, imaging guidelines, and scanner access may change over time	Required before using the model in live clinical workflow
External validation	Test across different hospitals, departments, specialties, and EHR/radiology systems	Repeat imaging patterns are highly dependent on institutional practice norms	Supports generalizability and identifies site-specific retraining needs
Subgroup and fairness assessment	Evaluate performance by age, sex, language, insurance status, race/ethnicity where appropriate and permitted, care setting, specialty, and chronic disease group	The model could over-flag patients with fragmented care or incomplete documentation	Fairness review should focus on unequal alert burden and unequal preservation of appropriate follow-up
Clinical usefulness	Decision-curve analysis, alert yield, appropriateness-review yield, avoided unsupported repeat orders	Predictive accuracy alone does not prove operational value	Deployment should be justified only if alerts improve review quality without delaying needed imaging
Interpretability	Temporal contribution summaries, attention review where valid, feature attribution, explanation of recent comparable exams and documented recommendations	Clinicians need to understand why the patient is predicted to receive repeat imaging	Explanations should show recognizable clinical drivers, not opaque model confidence alone
Alert-fatigue monitoring	Alert frequency, override rate, dismissal reasons, time added to order entry, clinician feedback	Even accurate models can fail if they interrupt urgent or appropriate care too often	Thresholds should be governed locally and adjusted to avoid excessive low-value nudges
Safety and appropriateness review	Manual review of sampled alerts, false-positive analysis, false-negative review, assessment of delayed or missed indicated imaging	The model must not function as a blunt imaging-reduction tool	Governance should emphasize appropriate imaging, not simple volume reduction
Data quality surveillance	Monitor missing outside imaging, incomplete indications, NLP extraction errors, delayed report availability, documentation variation	Missing or uneven data could make justified repeat imaging appear unsupported	Data-quality warnings should accompany model outputs when key evidence is unavailable
Model drift monitoring	Track calibration, alert rates, repeat-order prevalence, modality mix, specialty mix, and override patterns over time	Imaging behavior and documentation standards may evolve after deployment	Drift surveillance should trigger review, recalibration, or retraining
Implementation evaluation	Pragmatic trial, stepped-wedge rollout, interrupted time-series	The manuscript proposes an operational intervention,	Final evaluation should measure appropriateness,

	analysis, or prospective pilot	not only a prediction algorithm	clinician trust, workflow burden, and unintended consequences
--	--------------------------------	---------------------------------	---

Limitations

Data quality and missing indications

The proposed model would depend on the quality, completeness, and timeliness of radiology information system and electronic health record data. Imaging codes may not fully capture clinical need, notes may omit important symptoms, specialist recommendations may be documented inconsistently, and prior imaging from outside institutions may be unavailable. Natural language processing studies show that report and note extraction can identify meaningful signals, but these methods remain sensitive to documentation style, terminology, and context [7, 8, 16, 17]. Consequently, incomplete or uneven documentation could cause the model to misinterpret clinically justified repeats as unsupported or miss important indications for follow-up.

Acceptability and alert fatigue

Clinical acceptability would be a major limitation because even accurate predictions may fail if they interrupt workflow, seem punitive, or provide explanations that clinicians consider irrelevant. Soft nudges can be ignored, while frequent alerts may contribute to fatigue and reduce attention to genuinely useful guidance. Prior decision support and imaging utilization studies suggest that implementation success depends on thoughtful design, local governance, transparency, and iterative refinement with ordering clinicians [4-6, 13]. The model should therefore be deployed cautiously as a learning system embedded in clinical practice rather than as a static rule engine.

Conclusion

A sequence learning model for predicting repeat diagnostic imaging orders would frame imaging utilization as a temporal and contextual prediction problem. By modelling prior imaging history, symptom documentation, specialist recommendations, ordering physician behavior, and recent test results, the system could anticipate repeat ordering before the final order-entry moment.

The key strength of this approach is its ability to preserve temporal structure while integrating heterogeneous clinical signals. Instead of relying only on static appropriateness rules, the model could identify how prior studies, evolving symptoms, clinician tendencies, and recent findings combine to create the conditions for repeat imaging.

Several challenges would remain before such a model could be used safely in practice. Data completeness, documentation variability, clinician acceptance, institutional culture, and the need to distinguish guideline-adherent follow-up from unwarranted variation would all shape whether the model improves care.

Real-world pilots within health systems would be needed to assess operational feasibility, clinician trust, and effects on low-value imaging. The most useful implementation would not simply reduce imaging volume, but would support better-timed, better-documented, and more appropriate diagnostic imaging decisions.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 26 Aug 2024

Revised: 23 Sep 2024

Accepted: 06 Nov 2024

Published online: 25 February 2025

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

- Kjelle E, Andersen ER, Soril LJ, van Bodegom-Vos L, Hofmann BM. Interventions to reduce low-value imaging: a systematic review of interventions and outcomes. *BMC Health Serv Res.* 2021;21(1):983.
- Kjelle E, Andersen ER, Krokeide AM, Soril LJ, van Bodegom-Vos L, Clement FM, et al. Characterizing and quantifying low-value diagnostic imaging internationally: a scoping review. *BMC Med Imaging.* 2022;22(1):73.
- Kjelle E, Brandsæter IØ, Andersen ER, Hofmann B. Sustainability in healthcare by reducing low-value imaging: a narrative review. *Radiography.* 2024;30:30-4.
- Ip IK, Lacson R, Hentel K, Malhotra S, Darer J. Predictors of provider response to clinical decision support: lessons learned from the Medicare Imaging Demonstration. *AJR Am J Roentgenol.* 2017;208(2):351-7.
- Palen TE, Sharpe RE Jr, Shetterly SM, Steiner JF. Randomized clinical trial of a clinical decision support tool for improving the appropriateness scores for ordering imaging studies in primary and specialty care ambulatory clinics. *AJR Am J Roentgenol.* 2019;213(5):1015-20.
- Rezaii PG, Fredericks N, Lincoln CM, Hom J, Willis M, Burleson J, et al. Assessment of the radiology support, communication and alignment network to reduce medical imaging overutilization: a multipractice cohort study. *J Am Coll Radiol.* 2020;17(5):597-605.
- Brown AD, Marotta TR. A natural language processing-based model to automate MRI brain protocol selection and prioritization. *Acad Radiol.* 2017;24(2):160-6.
- Carrodegua E, Lacson R, Swanson W, Khorasani R. Use of machine learning to identify follow-up recommendations in radiology reports. *J Am Coll Radiol.* 2019;16(3):336-43.

Sorin V, Barash Y, Konen E, Klang E. Deep learning for natural language processing in radiology: fundamentals and a systematic review. *J Am Coll Radiol.* 2020;17(5):639-48.

Shickel B, Tighe PJ, Bihorac A, Rashidi P. Deep EHR: a survey of recent advances in deep learning techniques for electronic health record (EHR) analysis. *IEEE J Biomed Health Inform.* 2017;22(5):1589-604.

Li Y, Rao S, Solares JR, Hassaine A, Ramakrishnan R, Canoy D, et al. BEHRT: transformer for electronic health records. *Sci Rep.* 2020;10(1):7155.

Rasmy L, Xiang Y, Xie Z, Tao C, Zhi D. Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ Digit Med.* 2021;4(1):86.

Brandsæter IØ, Andersen ER, Hofmann BM, Kjelle E. Drivers for low-value imaging: a qualitative study of stakeholders' perspectives in Norway. *BMC Health Serv Res.* 2023;23(1):295.

Lakhani P, Prater AB, Hutson RK, Andriole KP, Dreyer KJ, Morey J, et al. Machine learning in radiology: applications beyond image interpretation. *J Am Coll Radiol.* 2018;15(2):350-9.

Kalra A, Chakraborty A, Fine B, Reicher J. Machine learning for automation of radiology protocols for quality and efficiency improvement. *J Am Coll Radiol.* 2020;17(9):1149-58.

Heilbrun ME, Chapman BE, Narasimhan E, Patel N, Mowery D. Feasibility of natural language processing-assisted auditing of critical findings in chest radiology. *J Am Coll Radiol.* 2019;16(9):1299-304.

Nakamura Y, Hanaoka S, Nomura Y, Nakao T, Miki S, Watadani T, et al. Automatic detection of actionable radiology reports using bidirectional encoder representations from transformers. *BMC Med Inform Decis Mak.* 2021;21(1):262.

Valtchinov VI, Ip IK, Khorasani R, Schuur JD, Zurakowski D, Lee J, et al. Use of imaging in the emergency department: do individual physicians contribute to variation? *AJR Am J Roentgenol.* 2019;213(3):637-43.

Hofmann B, Andersen ER, Brandsæter IØ, Clement F, Elshaug AG, Bryan S, et al. Success factors for interventions to reduce low-value imaging: six crucial lessons learned from a practical case study in Norway. *Curr Probl Diagn Radiol.* 2024;53(6):670-6.

Li Y, Mamouei M, Salimi-Khorshidi G, Rao S, Hassaine A, Canoy D, et al. Hi-BEHRT: 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.

Yang Z, Mitra A, Liu W, Berlowitz D, Yu H. TransformEHR: transformer-based encoder-decoder generative model to enhance prediction of disease outcomes using electronic health records. *Nat Commun.* 2023;14(1):7857.

Kraljevic Z, Bean D, Shek A, Bendayan R, Hemingway H, Yeung JA, et al. Foresight: a generative pretrained transformer for modelling patient timelines using electronic health records: a retrospective modelling study. *Lancet Digit Health.* 2024;6(4):e281-e290.

Klang E, Barash Y, Soffer S, Bechler S, Resheff YS, Granot T, et al. Promoting head CT exams in the emergency department triage using a machine learning model. *Neuroradiology.* 2020;62(2):153-60.

Wong KA, Hatf A, Ryu JL, Nguyen XV, Makary MS, Prevedello LM. An artificial intelligence tool for clinical decision support and protocol selection for brain MRI. *AJNR Am J Neuroradiol.* 2023;44(1):11-6.

Lacson R, Eskian M, Licaros A, Kapoor N, Khorasani R. Machine learning model drift: predicting diagnostic imaging follow-up as a case example. *J Am Coll Radiol.* 2022;19(10):1162-9.

Shahbandegan A, Mago V, Alaref A, van der Pol CB, Savage DW. Developing a machine learning model to predict patient need for computed tomography imaging in the emergency department. *PLoS One*. 2022;17(12):e0278229.

Wamil M, Hassaine A, Rao S, Li Y, Mamouei M, Bidel Z, et al. Stratification of diabetes in the context of comorbidities, using representation learning and topological data analysis. *Sci Rep*. 2023;13(1):11478.

Antikainen E, Linnoosmaa J, Umer A, Oksala N, Eskola M, Van Gils M, et al. Transformers for cardiac patient mortality risk prediction from heterogeneous electronic health records. *Sci Rep*. 2023;13(1):3517.

Belavy DL, Tagliaferri SD, Buntine P, Saueressig T, Samanna C, Owen PJ, et al. Reducing low-value imaging for low back pain: systematic review with meta-analysis. *J Orthop Sports Phys Ther*. 2022;52(4):175-91.