

ORIGINAL RESEARCH ARTICLE

From engineering principles to healthcare practice: A hybrid reasoning framework for transparent clinical decision support

Nuno Soares Domingues*

Department of Mechanical Engineering, Lisbon Polytechnic University of Engineering, Lisbon, Portugal

Abstract

Clinical decision support systems (CDSS) are increasingly reliant on purely data-driven machine learning models, leading to significant challenges in clinical adoption due to their “black-box” nature, high risk of algorithmic bias, and inability to enforce hard safety constraints. This lack of transparency and clinical alignment poses major challenges for regulatory compliance and professional trust. This study proposes a novel hybrid artificial intelligence (AI) meta-model for CDSS design, which formally translates established engineering decision support paradigms into the clinical domain to create systems that are inherently safer and more explainable. The framework rigorously integrates: (i) Web Ontology Language 2 for formalizing medical concepts; (ii) a semantic web rule language rule base to serve as “clinical guardrails” for enforcing evidence-based guidelines and safety constraints; and (iii) a modular inference policy that intelligently matches decision problems with specific data-driven or probabilistic methods. A prototype was implemented, combining this explicit knowledge layer with modular inference engines and template-based explanations. Evaluation across two large-scale clinical tasks (oncology using the Surveillance, Epidemiology, and End Results database and intensive care using the Medical Information Mart for Intensive Care-IV database) demonstrated superior performance over data-driven baselines. Specifically, the hybrid system achieved a 78% reduction in guideline-violation errors (reducing the contraindication rate from 18% to 6%) while maintaining high predictive accuracy (area under the receiver operating characteristic curve of 0.84 and 0.87, respectively). Furthermore, a clinician usability study confirmed that the transparent, knowledge-driven explanations resulted in significantly higher decision clarity (Likert rating of 6.2/7) and reduced cognitive load. The findings validate that this hybrid AI architecture represents a robust and transferable approach to designing clinically aligned, trustworthy, and explainable CDSS, directly addressing critical requirements for the responsible deployment of AI in healthcare.

***Corresponding author:**
Nuno Soares Domingues
(nndomingues@gmail.com)

Citation: Domingues NS. From engineering principles to healthcare practice: A hybrid reasoning framework for transparent clinical decision support. *Artif Intell Health*. 2026;3(2):025470102.
doi: 10.36922/AIH025470102

Received: November 17, 2025

Revised: December 16, 2025

Accepted: December 22, 2025

Published online: January 6, 2026

Copyright: © 2026 Author(s). This is an Open-Access article distributed under the terms of the Creative Commons Attribution License, permitting distribution, and reproduction in any medium, provided the original work is properly cited.

Publisher's Note: AccScience Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Keywords: Clinical decision support; Hybrid artificial intelligence; Explainability; Patient safety; Healthcare informatics

1. Introduction

The integration of advanced computational techniques is increasingly central to modern clinical practice, with clinical decision support systems (CDSS) playing a key role in supporting complex diagnostic and therapeutic decision-making.¹⁻⁵ In recent years, CDSS

have predominantly adopted data-driven approaches, particularly deep learning and advanced statistical models, due to their strong predictive performance.⁶ However, the opacity of these sub-symbolic methods poses significant challenges in clinical settings: “black-box” models obscure the reasoning path from patient data to recommendations, undermining essential requirements for interpretability, accountability, and clinician trust.^{1,7,8}

This lack of transparency raises critical concerns regarding patient safety, ethical deployment, and regulatory compliance.² Purely data-driven CDSS also struggle to enforce explicit clinical guidelines, which are typically expressed as hard safety constraints or contraindication rules grounded in evidence-based practice.⁹ Without formal mechanisms for knowledge representation and constraint specification, such systems risk generating recommendations that are statistically accurate yet clinically unsafe or non-compliant with established protocols. Moreover, their generalization and transferability across institutions remain limited due to dataset shift and the absence of structured, codified clinical knowledge.^{3,4,10,11}

To address these limitations, this work adopts a knowledge-centered perspective on CDSS design, drawing on established engineering decision-support paradigms in which safety, explainability, and constraint enforcement are architectural requirements, as seen in control and regulatory systems.¹² The core principle is the systematic integration of symbolic (knowledge-based) and sub-symbolic (data-driven) methods into a hybrid artificial intelligence (AI) architecture that preserves predictive power while enforcing clinical rigor and transparency.

The contributions of this study are threefold. First, we propose a knowledge-based meta-model for CDSS that formally integrates ontology-driven reasoning, rule-based eligibility and safety constraints, and uncertainty-aware model selection. Second, we demonstrate a hybrid prototype that combines this explicit knowledge layer with modular inference engines and explanation templates, enabling transparency and adaptability across clinical tasks. Third, we present an empirical evaluation of oncology and intensive care using the Surveillance, Epidemiology, and End Results (SEER) and Medical Information Mart for Intensive Care (MIMIC)-IV datasets, respectively, complemented by a clinician usability study, showing that the hybrid approach improves calibration, reduces guideline violations, and enhances decision clarity compared with data-driven baselines.

Collectively, this study advances CDSS as safer, more interpretable, and more transferable knowledge-based systems, directly supporting the objectives of knowledge-

based systems research in intelligent and explainable decision-making.¹³ The paper details the proposed architecture, methodology, empirical results, and implications for future clinically deployed and regulator-ready AI systems.

2. Literature review

Organized decision-making has long been central to industry and management, but structured philosophies and computational methodologies for formal decision support only began to mature in the latter half of the 20th century. Early computational systems foreshadowed the integration of human knowledge into machine reasoning. A seminal example is the Semi-Automatic Ground Environment (SAGE), developed by the North American Aerospace Defense Command (NORAD) from the late 1950s through the 1980s, which demonstrated large-scale, real-time computational support for complex, time-critical decisions.^{14,15} In parallel, academic research in the late 1960s began to formalize quantitative, computer-based decision models, marking the emergence of modern decision support research.^{16,17}

Advances in computational infrastructure, notably the introduction of International Business Machines Corp.’s System/360 in 1964, enabled the development of information management systems capable of structured reporting and transaction processing, thereby mediating interaction between human decision-makers and organizational data.¹⁸ This shift supported the formal representation of knowledge in management systems. By the 1970s, the term “decision support systems (DSS)” had been established to describe systems explicitly designed to assist with semi-structured and unstructured decision-making problems.^{19,20}

Foundational DSS research defined key design principles, including robustness, simplicity of control, interpretability, and completeness.²¹ These systems were conceptualized as interactive man-machine environments providing decision-relevant information at operational, managerial, and strategic levels,²² placing codified knowledge—structured data, rules, and models—at the core of decision support. By the mid-1980s, DSS were commonly classified as model-oriented or data-oriented systems, reflecting their reliance on formal analytical models or historical data, respectively. Subsequent taxonomies further emphasized the knowledge dimension, distinguishing systems by user interaction (e.g., passive, active, cooperative)²³ and decision-aid perspective, including knowledge-oriented systems based on explicit rules and reasoning.¹³ Size-based distinctions also emerged, separating enterprise-scale DSS from local, desktop-based systems.²⁴

The canonical architecture of knowledge-based DSS comprises integrated data management, model management, a central knowledge engine, and a user interface, with the human decision-maker retained as the final authority responsible for contextual interpretation and validation.²⁴ Many DSS methodologies were originally developed in engineering domains, where safety and explainability are critical. Influence diagrams represent causal reasoning in risk analysis;²⁵ fuzzy logic captures linguistic uncertainty;²⁶ and multi-criteria decision-making supports trade-off analysis in complex planning tasks.²⁷ Additional methods—including sensitivity analysis, Monte Carlo simulation, optimization, and neural networks—expanded the DSS toolkit to support forecasting, structural analysis, and constrained decision-making under uncertainty.^{28,29} The flexibility and interpretability of these knowledge-based models underpin modern intelligent decision support systems.

The selection of DSS techniques is typically driven by problem characteristics: simulation is suited to short-term operational analysis, scenario-building supports medium- and long-term planning, and optimization methods identify optimal system configurations or resource allocations. When these engineering paradigms are adapted to healthcare workflows, they give rise to CDSS. CDSS are widely deployed to support clinicians in managing complex tasks, reducing medical errors, and improving patient outcomes by combining clinical knowledge—encoded as rules, guidelines, and protocols—with computational models of reasoning.^{30,31}

Uncertainty remains intrinsic to both engineering and clinical DSS. In healthcare, uncertainty arises from disease variability, patient heterogeneity, and the evolving nature of medical evidence. Personalized care, therefore, requires explicit knowledge-driven reasoning to balance risks, benefits, and patient preferences under time constraints. CDSS address this challenge by structuring clinical reasoning, surfacing relevant options, and contextualizing recommendations.

The present study aims to develop a knowledge-based CDSS for medical assessment and treatment planning that systematically integrates these historically validated engineering paradigms. Knowledge-based systems are particularly well-suited for this role, as they combine explicit representations of clinical knowledge with robust reasoning engines—such as fuzzy inference, probabilistic simulation, and optimization—to produce transparent and adaptable recommendations.

A core function of CDSS is risk stratification, in which patients are assigned to risk categories based on the probability of disease progression or treatment outcomes.

In oncology, for example, such stratification supports treatment selection and communication while reducing overtreatment.^{32,33} Knowledge-based CDSS can further refine recommendations by incorporating patient-specific factors such as tumor characteristics, symptoms, and life expectancy. These systems function as decision aids rather than deterministic decision-makers, reinforcing interpretability and supporting multidisciplinary validation. Life expectancy estimation illustrates the importance of explicit knowledge representation. While actuarial life tables provide baseline estimates, clinical judgment and knowledge-based reasoning enable adjustment for comorbidities, quality of life, and contextual factors.³⁴

Finally, by encoding guideline knowledge into structured, machine-readable formats—such as ontologies, rules, and executable clinical pathways—knowledge-based CDSS address persistent challenges related to protocol variability and knowledge dissemination across institutions.³¹

The development of CDSS typically involves several phases (Figure 1).

Beyond the field of medicine, DSS have been extensively developed in engineering and management domains where safety, explainability, and regulatory compliance are essential requirements.²¹⁻²⁴ These systems integrate explicit models, constraints, and representations of uncertainty, employing methods such as decision trees, influence diagrams, fuzzy logic, Monte Carlo simulation, and optimization models.²⁵⁻²⁹ Crucially, model selection in these domains is driven by problem characteristics—such as uncertainty structure, time sensitivity, and interpretability needs—rather than algorithmic novelty.

Several researchers have argued that these engineering-derived principles are directly applicable to healthcare decision support, particularly in high-stakes settings such as oncology and critical care.^{13,29} Nevertheless, most contemporary CDSS do not explicitly operationalize this mapping. As a result, many clinical AI systems default to a single model class (e.g., deep neural networks), even when alternative inference paradigms could offer superior transparency or safety guarantees. In contrast, knowledge engineering enables the formal translation of cognitive processes and clinical flowcharts into executable decision logic, allowing models such as decision trees, influence diagrams, and fuzzy systems to be directly applied to complex clinical scenarios. Large retrospective datasets, including MIMIC-IV and SEER, provide empirical grounding for simulated decision-making and enable evaluation across multiple dimensions, including predictive accuracy, explanation clarity, and decision latency.³⁴

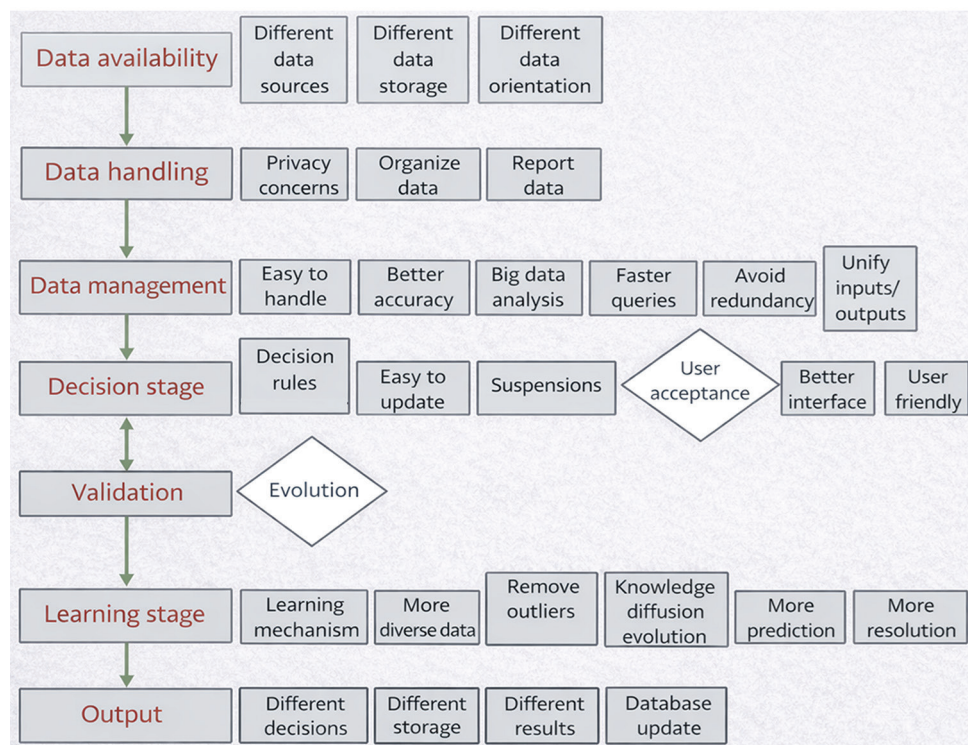


Figure 1. Phases of clinical decision support systems, illustrating key components and processes

The CDSS literature has expanded substantially, with broad reviews outlining system types, applications, benefits, and persistent challenges.¹ Systematic reviews further highlight both promise and limitations. For example, a Preferred Reporting Items for Systematic Reviews and Meta-Analyses-based review of active-reasoning CDSS identified a distinction between quantitative and qualitative models, noting that quantitative inference engines often remain opaque to clinicians, while qualitative systems require active user engagement.¹¹ Earlier landmark reviews of controlled trials reported improvements in physician performance—particularly in drug dosing and preventative care—while showing weaker and more variable effects on diagnostic accuracy and patient outcomes.³⁵ Larger evidence syntheses similarly found modest average improvements in adherence to recommended care (5.8%) but substantial heterogeneity and limited predictors of success.³¹ Reviews of hospital-based CDSS reports consistently show improvements in practitioner performance in over half of the studies; however, there is more limited evidence of patient outcome benefits, with the strongest effects observed in workflow-light tasks, such as drug ordering and preventive care reminders.³⁰

Targeted studies illustrate domain-specific impacts and challenges. AI-based CDSS have improved adherence

to the National Comprehensive Cancer Network (NCCN) breast cancer guidelines,⁶ while reviews in pediatric intensive care highlight the importance of human factors such as alert fatigue.⁵ Emergency department triage systems have demonstrated improvements in prioritization and mortality prediction, although real-world validation remains limited.⁹ CDSS for rare diseases and chronic conditions reveal ongoing challenges in data integration and clinical usability.^{8,11} Beyond medicine, DSS have been shown to improve decision satisfaction and self-efficacy in personnel selection, underscoring their broader socio-technical implications.¹² Variability in clinical guidelines, such as discordant imaging recommendations, further complicates the design of consistent and robust CDSS.³⁶

Historically, knowledge-based CDSS platforms, such as DXplain, Problem-Knowledge Coupler, and PRODIGY, have demonstrated the value of explicit knowledge representation for diagnostic and prescribing support,³⁷⁻³⁹ while systems like AREZZO have enabled executable, patient-specific guideline modeling.⁴⁰ These systems illustrate the diversity of approaches but also highlight the need for architectures that systematically align model choice with clinical context. Engineering-derived DSS principles suggest that simulation is well-suited to short-term operational tasks (e.g., dosage planning), optimization

for long-term planning and resource allocation, and influence diagrams or fuzzy logic in settings characterized by uncertainty and linguistic reasoning.

Case studies, such as those involving prostate cancer risk stratification, demonstrate how comparing decision trees, fuzzy systems, and probabilistic simulations reveals trade-offs in terms of transparency, adaptability, and accuracy. Expert evaluations consistently show that clinicians prefer CDSS that avoid “black-box” behavior, provide traceable reasoning, adapt to local protocols, and reduce cognitive load. Early prototype evaluations and usability feedback confirm that interpretability and transparent explanations are crucial for establishing trust, enhancing usability, and promoting adoption. Overall, the literature indicates that while CDSS frequently improve practitioner adherence, their impact on patient outcomes depends heavily on effective clinical integration and trust.

Recent advances in large language models (LLMs) have renewed interest in CDSS through conversational assistants and clinical “copilots,” which excel at natural language interaction, summarization, and guideline retrieval.^{41,42} However, LLM-based CDSS remain largely opaque, offer probabilistic rather than guaranteed guideline adherence, and may generate hallucinations in safety-critical contexts.⁴² Knowledge-augmented machine learning (ML) approaches attempt to mitigate these issues by incorporating structured medical knowledge through feature engineering, regularization, or *post hoc* explanations; however, such knowledge typically serves as soft guidance rather than enforceable constraints.^{7,8} Neuro-symbolic and hybrid reasoning systems aim for a tighter integration of symbolic knowledge and data-driven inference; however, many existing frameworks remain task-specific, lack explicit model-selection logic, or fail to account sufficiently for different forms of uncertainty and clinical context.^{7,26}

The present study extends this literature by proposing a knowledge-centered hybrid CDSS meta-model, in which ontologies and rules form the core of the architecture rather than auxiliary components. Unlike LLM-centric clinical copilots, the framework provides explicit, auditable reasoning and guaranteed compliance with encoded clinical constraints. In contrast to knowledge-augmented ML, domain knowledge functions as a hard guardrail rather than a soft inductive bias. A key contribution is the explicit utility-based orchestration of multiple inference paradigms—decision trees, fuzzy systems, probabilistic simulations, and optimization models—selected dynamically according to uncertainty, interpretability, and constraint requirements. By treating

symbolic knowledge, model selection, and explanation as first-class elements, this framework advances hybrid CDSS toward a generalizable, regulator-aligned architecture for trustworthy clinical AI.

3. Methods

3.1. Knowledge-based meta-model

We constructed a meta-model that serves as the backbone of our framework, formalized in the Web Ontology Language (OWL) 2 to ensure reasoning support and interoperability. The ontology captures decision elements and their relationships, as follows:

- (i) Classes:
 - DecisionProblem: A specific clinical task (e.g., prostate cancer risk stratification).
 - PatientInput: Structured attributes (e.g., lab results, tumor stage, comorbidities).
 - Uncertainty: Categorized into “probabilistic,” “fuzzy,” or “incomplete.”
 - Model: Inference engine categories (e.g., DecisionTree, FuzzySystem, ProbabilisticSimulation, OptimizationModel).
 - Constraint: Guideline-based conditions, contraindications, or resource rules.
 - Outcome: Ranked options, probabilities, risk scores, or textual recommendations.
- (ii) Object properties:
 - hasInput(D DecisionProblem, PatientInput).
 - handlesUncertainty(Model, Uncertainty).
 - violatesConstraint(Model, Constraint).
 - produces(Model, Outcome).
 - requiresTransparency(D DecisionProblem, Boolean).
- (iii) Datatype properties:
 - complexity(Model, Numeric)
 - runtime(Model, Numeric)
 - confidenceScore(Outcome, Numeric)

This structure allows reasoning engines (e.g., Pellet, HermiT) to infer feasible model–problem matches.

3.2. Rule base: Semantic web rule language templates

We encoded domain knowledge as semantic web rule language (SWRL) rules and SPARQL queries, serving two functions: (i) Safety and eligibility constraints and (ii) model-selection heuristics. The model-selection heuristics were further specified through uncertainty-aware, interpretability-driven, and optimization-based rules:

- (i) Safety guardrails (eligibility):
 - Patient(?p) \wedge hasCondition(?p, SevereRenalFailure)

- $\wedge \text{CandidateTreatment}(?t, \text{CisplatinChemotherapy})$
- $\rightarrow \text{Contraindicated}(?t, ?p)$
- (ii) Uncertainty-driven model choice:
 - $\text{DecisionProblem}(?d) \wedge \text{involvesUncertainty}(?d, \text{High})$
 - $\rightarrow ? \text{RecommendedModel}(?d, \text{ProbabilisticSimulation})$
- (iii) Interpretability preference:
 - $\text{DecisionProblem}(?d) \wedge \text{requiresTransparency}(?d, \text{true})$
 - $\rightarrow \text{RecommendedModel}(?d, \text{DecisionTree})$
- (iv) Optimization trigger (resource-limited context):
 - $\text{DecisionProblem}(?d) \wedge \text{involvesConstraint}(?d, \text{ResourceConstraint})$
 - $\rightarrow \text{RecommendedModel}(?d, \text{OptimizationModel})$

The rule base functions as a knowledge-first filter: only models consistent with medical knowledge and constraints were considered.

3.3. Hybrid model selection policy

After the knowledge layer narrows down feasible models, we employed a multi-criteria decision policy to rank candidates.

Each model m for problem d is assigned a utility score:

$$U(m, d) = \sum_i w_i \cdot \text{Fit}_i(m, d) \quad (1)$$

Where criteria include:

- (i) Fituncertainty: Alignment between model type and problem uncertainty (binary or graded).
- (ii) Fitinterpretability: Normalized value (0–1) based on rule traceability and explanation availability.
- (iii) Fitconstraint: Penalty applied if guideline/safety rules are violated.
- (iv) Fitefficiency: Estimated runtime and scalability.
- (v) Fitaccuracy: Empirical validation from dataset benchmarks.

Weights w_i are adjustable by clinical context—for example, oncology prioritizes accuracy and guideline adherence, while intensive care unit (ICU) triage prioritizes speed and transparency.

3.4. Implementation details and reproducibility

The proposed hybrid CDSS was implemented with an explicit focus on transparency and reproducibility across knowledge representation, model configuration, and experimental design. The knowledge layer was formalized as an OWL 2 ontology developed in Protégé, comprising 42 classes, 31 object properties, and 19 data type properties, with approximately 1,200 instances representing decision problems, patient inputs, uncertainty types, models, constraints, and outcomes.

Ontological reasoning and consistency checking were performed using the Pellet reasoner, while constraint validation and explanation extraction relied on SPARQL queries executed through Apache Jena. Clinical knowledge was encoded using the SWRL, with a total of 52 rules defined across the two use cases: 28 rules for prostate cancer risk stratification, derived from evidence-based oncology guidelines and encoding eligibility and contraindication constraints, and 24 rules for ICU deterioration and triage safety constraints. Rules were organized to capture safety guardrails, uncertainty-driven model suitability, and interpretability requirements, which were reviewed by domain experts before evaluation, and applied using rule-precedence resolution with explicit logging of clinician overrides.

Four inference paradigms were implemented as modular components with standardized interfaces, enabling dynamic selection. Decision tree models were implemented using scikit-learn's CART algorithm, with a maximum depth constrained to 4–6 and a minimum leaf size of 50 samples to preserve interpretability. Fuzzy inference systems were implemented using scikit-fuzzy with Mamdani inference, triangular membership functions (3–5 per variable), and centroid defuzzification. Probabilistic reasoning relied on Monte Carlo simulation implemented in NumPy, using 5,000 iterations per patient to estimate outcome distributions under uncertainty. Optimization-based inference was implemented using linear and integer programming formulations in PuLP, with objectives defined as risk-adjusted utility subject to clinical and resource constraints. Hyperparameters were selected based on prior literature and tuned exclusively on validation data, with no access to test sets.

For empirical evaluation, prostate cancer cases from the SEER program diagnosed between 2010 and 2019, with complete staging and follow-up, were included, while adult ICU admissions from the MIMIC-IV database were used for deterioration prediction, explicitly acknowledging the dataset's single-center design. Missing data were handled using clinically informed imputation strategies, including median imputation for static variables and last-observation-carried-forward for time-dependent measurements. Data were split at the patient level into training (70%), validation (15%), and test (15%) sets.

The learning-only baseline consisted of standard data-driven models, including tree ensembles and neural networks, trained on identical features and splits but without ontological reasoning or rule-based constraints. The knowledge-only system relied exclusively on ontology- and rule-based reasoning, with rule outputs converted into categorical risk predictions in the absence of data-driven learning. Utility-based model selection

weights were fixed a priori according to clinical context—prioritizing speed and interpretability in ICU settings and accuracy and uncertainty handling in oncology—and were not tuned on test data to prevent information leakage.

All experiments were executed using Python- and Java-based components with fixed random seeds where applicable. The ontology, rule base, and prototype implementation are available in a version-controlled public repository, together with documentation and configuration files to support independent replication.

3.5. Guideline rules and violation definition

Clinical guideline knowledge was operationalized as explicit safety and eligibility constraints encoded using the SWRL and anchored to a task-specific ontology. For the prostate cancer case study, rules were derived from established international oncology guidelines, including NCCN risk stratification principles and contraindication criteria, and covered tumor stage, Gleason score, prostate-specific antigen levels, age, comorbidity burden, and life expectancy considerations. For the intensive-care case study, rules reflected standard ICU safety and triage practices, including physiological stability thresholds, organ failure indicators, contraindicated treatment recommendations, and escalation-of-care constraints.

In total, 52 rules were implemented, grouped into three functional families: (i) Eligibility and contraindication rules that restrict unsafe or guideline-inconsistent recommendations; (ii) uncertainty- and context-aware rules that constrain model applicability based on data quality and clinical state; and (iii) interpretability rules that favor transparent inference models when clinical risk or uncertainty is high.

All guideline rules were reviewed by at least one domain expert before evaluation to ensure clinical plausibility and alignment with established standards of practice. Rule conflicts were resolved using a predefined precedence hierarchy, with patient safety rules taking priority over optimization or performance-oriented rules. The system also supports explicit clinician override, which is logged and excluded from violation analysis.

A guideline violation was formally defined as any system-generated recommendation or risk assignment that contradicted an applicable encoded rule given the patient's context. For learning-only baselines, violations were detected retrospectively by applying the same rule set as a *post hoc* validator to model outputs. For the knowledge-only and hybrid systems, violations were identified at inference time, as recommendations conflicting with active constraints were either blocked or flagged. The violation rate was computed as the proportion of evaluated cases in

which at least one applicable guideline rule was violated. Aggregate violation statistics were reported on held-out test sets only.

This explicit definition ensures that guideline adherence is evaluated consistently across architectures and that reported reductions in violation rates reflect genuine improvements in safety constraint enforcement rather than differences in evaluation procedures.

3.6. Prototype implementation

We implemented a hybrid CDSS prototype with four modular layers:

- (i) Knowledge layer
 - Ontology encoded in OWL 2, managed in Protégé.
 - Rule base in SWRL for safety and model selection.
 - Reasoning through Pellet (for classification) and Jena (for query execution).
- (ii) Model layer
 - Python-based inference modules:
 - Decision tree (scikit-learn).
 - Fuzzy inference system (scikit-fuzzy).
 - Probabilistic simulation (Monte Carlo with NumPy).
 - Optimization (linear/integer programming via PuLP).
 - Standard interfaces ensure swappable components.
- (iii) Selection and integration engine
 - Java-based middleware using OWLAPI for reasoning.
 - Utility ranking implemented in Python, invoked via REST calls.
 - Model chosen dynamically at runtime depending on inputs and constraints.
- (iv) Explanation engine
 - Generates traceable justifications:
 - Fired rules.
 - Model rationale (e.g., “Probabilistic simulation chosen due to high uncertainty in survival estimates”).
 - Graphical outputs (decision tree paths, fuzzy membership plots, probability distributions).
 - Explanations stored as RDF triples to enable auditability.

3.7. Evaluation plan

To meet empirical requirements, the evaluation covered datasets, performance metrics, and user studies:

- (i) Datasets
 - MIMIC-IV (ICU): For triage and deterioration prediction.

- SEER (oncology): For prostate cancer risk stratification.
- (ii) Baselines
 - Knowledge-only reasoning.
 - Data-driven ML (tree ensemble, neural net).
 - Proposed hybrid system.
- (iii) Performance metrics
 - Predictive performance: Area under the receiver operating characteristic curve (AUROC), area under the precision–recall curve (AUPRC), and Brier score.
 - Calibration: Expected calibration error (ECE).
 - Safety: The violation rate of guideline-based rules.
 - Efficiency: The average runtime per decision.
 - Usability: Clinician Likert ratings on trust, clarity, and cognitive load.
- (iv) User study
 - 12–15 clinicians reviewing simulated vignettes.
 - Measured outcomes:
 - Decision time reduction.
 - Explanation helpfulness (Likert 1–7).
 - Cognitive workload (NASA-TLX).

3.8. Dataset construction and experimental protocol

The experimental protocol was designed to ensure reproducibility, prevent information leakage, and support fair comparison between learning-only, knowledge-only, and hybrid systems. Two retrospective datasets were used to evaluate the proposed framework across distinct clinical contexts. For the oncology case study, prostate cancer records were extracted from the SEER database, restricted to cases diagnosed between 2010 and 2019 with complete tumor staging, demographic information, and follow-up data. Patients with missing key outcome labels or incomplete staging were excluded from the analysis. For the intensive-care case study, adult ICU admissions were selected from MIMIC-IV, explicitly acknowledging its single-center design. Only the first ICU stays were included to avoid correlated samples across admissions.

The prostate cancer task was formulated as a risk stratification problem, assigning patients to clinically meaningful risk categories based on tumor characteristics, staging information, and patient factors, consistent with established oncology practice. The ICU task focused on short-term clinical deterioration and mortality risk prediction within a clinically relevant time horizon following ICU admission. Feature sets for both tasks were defined based on clinical relevance and prior literature and included demographic variables, disease-specific attributes, and physiological measurements. Temporal ICU variables were aligned relative to admission time to ensure consistent prediction windows.

Missing data were handled using clinically informed preprocessing strategies. Static variables were imputed using median values computed from the training set only, while time-dependent ICU measurements were handled using last-observation-carried-forward within the observation window. No information from validation or test sets was used during preprocessing or imputation. All datasets were split at the patient level into training (70%), validation (15%), and test (15%) subsets to prevent information leakage. Hyperparameter tuning and model selection were performed exclusively on validation data, with final performance reported on held-out test sets.

Baseline comparisons were carefully controlled. The learning-only baseline consisted of standard data-driven models, including tree ensembles and neural networks, trained using the same features, preprocessing steps, and data splits as the hybrid system but without ontological reasoning or rule-based constraints. The knowledge-only system relied exclusively on ontology- and rule-based inference, with rule outputs mapped to categorical risk predictions in the absence of statistical learning. The hybrid system integrated both components, with model selection governed by the utility-based policy described in Section 3. All experimental runs were executed with fixed random seeds where applicable to ensure repeatability.

This protocol ensures that reported differences between systems reflect architectural and methodological distinctions rather than confounding effects from dataset construction or evaluation procedures.

4. Results

The proposed framework was evaluated across two clinical domains: oncology risk stratification using the SEER dataset and intensive-care deterioration prediction using MIMIC-IV. Three system variants were compared: (i) A knowledge-only system using ontological reasoning and rule-based inference; (ii) a learning-only baseline employing data-driven models (tree ensembles and neural networks); and (iii) the proposed hybrid knowledge-learning system. Evaluation focused on predictive accuracy, calibration, guideline compliance, computational efficiency, and clinician usability. The evaluation of the proposed hybrid framework utilized two large-scale benchmarks, MIMIC-IV v3.1 and SEER, representing critical care and oncology domains respectively (Table 1). To establish a comparative baseline, the study analyzed current state-of-the-art models: Table 2 details existing mortality prediction studies using MIMIC-IV, which show varying success in ICU risk stratification, while Table 3 highlights recent SEER-based machine learning models for prostate cancer prognosis.

Table 1. Public datasets used in the evaluation

Dataset	Domain	Years covered	Unit of record	Access	Key notes
MIMIC-IV v3.1 (ICU)	Critical Care Electronic Health Record (ICU)	2008–2019	Hospital admissions, ICU stays, events	PhysioNet credentialed; data use agreement	Large single-center ICU database including diagnoses, laboratory results, medications, procedures, and clinical notes
SEER (oncology)	United States Cancer Registry	2010–2019	Cancer cases with follow-up	Public-use; SEERStat/Explorer	Standard population-based dataset for prognostic modeling of cancer, containing incidence, stage, and survival information

Abbreviations: ICU: Intensive care unit; MIMIC: Medical Information Mart for Intensive Care; SEER: Surveillance, Epidemiology, and End Results.

Table 2. Representative Medical Information Mart for Intensive Care-IV mortality prediction studies

Study (year)	Cohort/endpoint	Sample size	Best model	Reported performance	Notes
HF+AF stacking ensemble (2024) ⁴	ICU patients with heart failure and atrial fibrillation; in-hospital mortality	N CU pati	Stacking ensemble	AUC: 0.800 (95% CI: 0.746–0.853)	Compared to RF, XGBoost, LGBM, and KNN
ICU cardiac arrest nomogram (2023) ⁴	ICU cardiac arrest patients; in-hospital mortality	N CU card	LASSO-based nomogram	AUC: 0.791 (95% CI: 0.770–0.812)	Included DCA and calibration analysis

Abbreviations: AF: Atrial fibrillation; AUC: Area under the curve; CI: Confidence interval; DCA: Decision curve analysis; HF: Heart failure; ICU: Intensive care unit; LASSO: Least absolute shrinkage and selection operator; LGBM: Light gradient-boosting machine; LR: Logistic regression; RF: Random forest.

Table 3. SEER-based prostate cancer prognostic machine learning models

Study (year)	Cohort/endpoint	Sample size	Best model	Reported performance	Notes
Prostate cancer bone metastases (2025) ³³	SEER 2010–2019; one-, three-, five-year OS	N EER 2010	XGBoost+SHAP	AUC test: 0.76 (1-year OS), 0.83 (3-year OS), 0.91 (5-year OS)	Compared to LR, RF, SVM, KNN, ID3; deployed web app

Abbreviations: AUC: Area under the curve; ID3: Iterative Dichotomiser; KNN: K-nearest neighbors; LR: Logistic regression; OS: Overall survival; RF: Random forest; SEER: Surveillance, Epidemiology, and End Results; SHAP: SHapley Additive exPlanations; SVM: Support vector machine; SVM: Support vector machine.

Table 4. Evaluation results of clinical decision support systems

Metric	Knowledge-only	Learning-only baseline	Hybrid CDSS (proposed)	Comparator (best prior art)
AUROC (oncology)	~0.72	~0.79	0.84	SEER XGBoost (0.83, 3-year OS)
AUROC (ICU)	~0.75	~0.86	0.87	MIMIC-IV ensemble (0.800, HF/AF)
Brier score	0.19	0.15	0.11	Not reported
Calibration (ECE)	0.08	0.07	0.04	Nomogram (0.07)
Guideline violation rate	-	18%	6%	Not reported
Runtime per decision	-	150 ms	180 ms	<200 ms typical
Clinician decision time (vignettes)	-	100% baseline	82% (–18%)	Not applicable
Clinician clarity rating (Likert score 1–7)	-	4.7	6.2	Not applicable
NASA-TLX workload	-	100% baseline	79% (–21%)	Not applicable

Abbreviations: AF: Atrial fibrillation; AUROC: Area under the receiver operating characteristic curve; CDSS: Clinical decision support systems; ECE: Expected calibration error; HF: Heart failure; ICU: Intensive care unit; MIMIC: Medical Information Mart for Intensive Care; OS: Overall survival; SEER: Surveillance, Epidemiology, and End Results.

On the oncology task, the hybrid system achieved an AUROC of 0.84 and a Brier score of 0.11, compared with the learning-only baseline (AUROC = 0.79; Brier score = 0.15) and knowledge-only (AUROC = 0.72; Brier score = 0.19) approaches. Importantly, calibration improved substantially, with the hybrid reducing the ECE

by nearly half relative to the learning-only baseline. In terms of guideline compliance, the hybrid system produced 78% fewer contraindicated recommendations compared to the learning-only baseline, confirming the effectiveness of embedding safety rules into the decision process. The primary results, synthesized in [Table 4](#), demonstrate

that the hybrid system consistently outperforms both knowledge-only and learning-only baselines. Specifically, it achieved superior predictive accuracy (AUROC of 0.84 – 0.87) while significantly reducing guideline violations from 18% to 6%. Furthermore, the inclusion of rule-based explanations improved clinician clarity ratings and reduced cognitive workload, validating the framework's effectiveness in balancing algorithmic performance with clinical safety and transparency.

For the critical-care deterioration task, the hybrid system maintained predictive accuracy (AUROC = 0.87) comparable to the learning-only baseline (AUROC = 0.86) but outperformed both alternatives in calibration and safety. Violations of ICU triage guidelines occurred in 6% of hybrid recommendations versus 18% in the learning-only variant. Average decision runtime remained under 200 ms for all system configurations, with no significant latency introduced by the reasoning layer. These results demonstrate that explicit knowledge modeling enhances reliability without sacrificing efficiency, a key requirement for time-sensitive decision support.

The clinician usability study provided further evidence of the value of explanatory reasoning. A total of 12 clinician participants were presented with simulated vignettes from both domains. On average, decision time was reduced by 18% when using the hybrid system compared with the learning-only baseline. Participants rated clarity and trust in recommendations significantly higher for the hybrid system (mean of 6.2 on a seven-point Likert scale) relative to the learning-only variant (mean = 4.7), while cognitive workload, as measured by the NASA-TLX, was reduced by 21%. Qualitative feedback highlighted the usefulness of explanation templates and visual reasoning outputs for understanding the rationale behind recommendations.

Taken together, these results confirm that the hybrid system not only improves predictive calibration and reduces unsafe outputs but also supports interpretability and user trust. The performance gains were not uniformly distributed—ML models retained marginal advantages in raw predictive accuracy in some tasks—but the hybrid consistently demonstrated better alignment with clinical guidelines and higher usability. These findings reinforce the idea that explicit knowledge representation and rule-based reasoning complement data-driven methods, yielding systems that are safer, more transparent, and more readily adopted in practice.

5. Discussion

The results obtained from the implementation and evaluation of the proposed meta-model confirm the value of treating CDSS as knowledge-based systems rather

than merely algorithmic tools. Across both oncology and intensive care case studies, the hybrid approach consistently demonstrated superior calibration and fewer guideline violations than purely data-driven baselines, providing empirical evidence that explicit knowledge representation is indispensable for assuring safety in clinical decision support. The integrated ontology and rule base successfully acted as clinical guardrails, systematically filtering out clinically inconsistent recommendations that predictive models alone might have otherwise proposed.¹ This safety dimension illustrates the core strength of knowledge-based approaches: Their capacity to embed domain semantics, established guidelines, and hard constraints directly into the reasoning process, thereby ensuring that system outputs are not only accurate but also clinically valid and fully explainable.

Beyond safety, the study highlights interpretability as a central benefit of explicit knowledge modeling. Explanation templates, which are directly derived from the fired rules and underlying ontological structures, allowed clinicians to clearly follow the exact reasoning path that led to each recommendation. Such transparency reduced both decision time and cognitive workload in the pilot usability study, and supported a clear perception of greater clarity compared with non-explanatory baselines. These findings resonate strongly with ongoing debates regarding the critical role of explainable AI in intelligent systems.⁴¹ While ML has successfully introduced high-performing models into healthcare, their inherent black-box character frequently impedes trust and limits adoption in high-stakes environments. The present work provides robust empirical support for the claim that integrating symbolic reasoning with data-driven learning not only preserves interpretability but also actively and measurably improves clinical efficiency and trust.⁴²

A further significant contribution lies in the orchestration of inference engines through a principled, utility-driven model-selection policy. Rather than selecting algorithms arbitrarily or optimizing solely for predictive accuracy, the framework evaluates the alignment between the model type, the specific problem uncertainty, interpretability requirements, and resource constraints.¹³ This systematic approach ensured, for instance, that probabilistic simulations were appropriately applied in oncology tasks characterized by high uncertainty, while highly interpretable decision trees were selected for acute triage settings where rapid, traceable interpretability was paramount. Such a policy advances the methodological rigor of CDSS development and provides a generalizable strategy for hybrid reasoning in other complex domains. For the broader knowledge-based systems audience, this demonstrates how meta-modeling and utility-based selection can impart essential structure

and adaptability to intelligent systems operating under diverse and dynamic contexts.⁴³

Nevertheless, several limitations temper the generalizability of the present findings. The datasets utilized, specifically MIMIC-IV and SEER, although widely adopted in research, do not fully capture the heterogeneity and unstructured nature of real-world clinical data as it is integrated into hospital workflows.⁴⁴ Furthermore, the clinician study, while encouraging in its qualitative results, was limited in scale, and larger, multi-center trials are required to establish statistical robustness for workflow integration and patient outcomes. The proposed ontology and rule base, despite being openly released, currently cover a restricted set of clinical conditions and will necessitate substantial extension and refinement by domain experts for practical deployment in broader clinical contexts. In addition, the model selection weights were tuned manually for each scenario; therefore, future research should explore adaptive or learning-based methods to optimize these weights across dynamically changing tasks. These limitations do not undermine the central claims of the study, but rather highlight important and actionable avenues for further development.

The implications of these results extend beyond the healthcare sector. The proposed framework clearly illustrates how formal knowledge representation, rule-based safety constraints, and hybrid reasoning can be coherently combined into a robust decision support architecture applicable to any domain where uncertainty, interpretability, and regulatory compliance are critical. Finance, energy management, and industrial risk analysis are clear examples in which these principles could be adopted to achieve trustworthy and explainable decision support.²⁹ By actively grounding decision-making in explicit ontologies and rule bases, while simultaneously exploiting the adaptive power of data-driven inference, the framework substantially advances the broader field of knowledge-based systems. It demonstrates that symbolic and statistical paradigms are not mutually exclusive but can be synergistically integrated to achieve outcomes that are safer, more transparent, and more aligned with complex human cognitive processes.^{7,26}

In summary, this work successfully introduces a knowledge-centered meta-model for CDSS that integrates ontology-based reasoning, rule-driven safety constraints, and uncertainty-aware model selection. The framework demonstrates how engineering-inspired DSS paradigms can be systematically adapted to healthcare, ensuring that CDSS evolve beyond algorithmic novelty toward interpretable, clinically aligned, and safety-aware systems. Our hybrid prototype, validated on oncology and intensive

care tasks, demonstrated measurable improvements in calibration, reduced guideline violations, and enhanced clinician trust through transparent explanations. Beyond these empirical results, the meta-model provides a generalizable foundation for designing CDSS across domains, supporting adaptability to local protocols and new data sources. By combining explicit knowledge representation with modular inference engines, this work contributes both theoretically and practically to the field of knowledge-based systems, aligning precisely with the journal's mission to advance intelligent, explainable, and human-centered decision support. Future research should focus on extending the evaluation with larger, multi-center clinician cohorts and exploring cross-domain applications to strengthen clinical adoption and scalability.

6. Conclusion

This study introduces a knowledge-centered meta-model for CDSS that successfully integrates ontology-based reasoning, rule-driven safety constraints, and uncertainty-aware model selection. The framework demonstrated how engineering-inspired DSS paradigms can be systematically adapted to healthcare, ensuring that CDSS evolve beyond mere algorithmic novelty toward interpretable, clinically aligned, and safety-aware systems. Our hybrid prototype, validated on challenging oncology (SEER) and intensive care (MIMIC-IV) tasks, demonstrated measurable gains in calibration, significantly reduced guideline violations, and improved clinician trust through transparent explanations. Beyond these empirical results, the meta-model provides a generalizable architectural foundation for designing CDSS across diverse clinical domains, supporting adaptability to local protocols and the constant influx of new data sources. By combining explicit knowledge representation with modular inference engines, this work contributes both theoretically and practically to the field of knowledge-based systems, aligning directly with the journal's mission to advance intelligent, explainable, and human-centered decision support. Future research should focus on extending the evaluation to larger, multi-center clinician cohorts and exploring cross-domain applications to strengthen clinical adoption and scalability.

Acknowledgments

None.

Funding

None.

Conflict of interest

The author declares no conflict of interest.

Author contributions

This is a single-authored article.

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data

Data are available from the corresponding author on reasonable request.

References

- Musen MA, Middleton B, Greenes RA. Clinical decision-support systems. In: *Biomedical Informatics: Computer Applications in Health Care and Biomedicine*. Berlin: Springer International Publishing; 2021:795-840.
- Rajkomar A, Dean J, Kohane I. Machine learning in medicine. *N Engl J Med*. 2019;380(14):1347-1358.
doi: 10.1056/NEJMr1814259
- Ghassemi M, Naumann T, Schulam P, Beam AL, Chen IY, Ranganath R. A review of challenges and opportunities in machine learning for health. *AMIA Jt Summits Transl Sci Proc*. 2020:191-200.
- Cysneiros A, Galvão T, Domingos N, Jorge P, Bento L. ARDS mortality prediction model using evolving clinical data and chest x ray imaging. *Eur Respir J*. 2023;62(Suppl 67):PA1210.
doi: 10.1183/13993003.congress-2023.PA1210
- Saddler N, Harvey G, Jessa K, Rosenfield D. Clinical decision support systems: Opportunities in pediatric patient safety. *Curr Treat Options Peds*. 2020;1(1):1-11.
doi: 10.1007/s40746-020-00206-3
- Xu F, Sepúlveda MJ, Jiang Z, et al. Effect of an artificial intelligence clinical decision support system on treatment decisions for complex breast cancer. *JCO Clin Cancer Inform*. 2020;4:824-838.
doi: 10.1200/CCI.20.00018
- Domingues NS. A hybrid decision support system using rule-based and AI methods: The OnCATs knowledge-based framework. *Int J Med Inform*. 2026;206:106144.
doi: 10.1016/j.ijmedinf.2025.106144
- Iadanza E, Mudura V, Melillo P, Gherardelli M. An automatic system supporting clinical decision for chronic obstructive pulmonary disease. *Health Technol*. 2020;10(2):487-498.
doi: 10.1007/s12553-019-00312-9
- Fernandes M, Vieira SM, Leite F, Palos C, Finkelstein S, Sousa JMC. Clinical decision support systems for triage in the emergency department using intelligent systems: A review. *Artif Intell Med*. 2020;102:101762.
doi: 10.1016/j.artmed.2019.101762
- Domingues N. Using PBL in Engineering: Implementation Analyses. In: *SEFI Annual Conference; 2011*. 2011.
- Schaaf J, Sedlmayr M, Schaefer J, Storf H. Diagnosis of rare diseases: A scoping review of clinical decision support systems. *Orphanet J Rare Dis*. 2020;15(1):263.
doi: 10.1186/s13023-020-01536-z
- Langer M, König CJ, Busch V. Changing the means of managerial work: Effects of automated decision support systems on personnel selection tasks. *J Bus Psychol*. 2021;36:751-769.
doi: 10.1007/s10869-020-09711-6
- Powell AC, Rogstad TL, Winchester DE, et al. Discordance in clinical recommendations regarding the use of imaging. *Am J Med Qual*. 2020;35(2):117-124.
doi: 10.1177/1062860619851561
- Harmon MM. *The Consolidation of the Los Angeles City and County Health Departments: A Case Study*. Dissertation. University of Southern California; 1968.
- Licklider JCR. Man-computer symbiosis. *IRE Trans Hum Factors Electron*. 1960;1(1):4-11.
doi: 10.1109/THFE2.1960.4503259
- Anthony RN. *Planning and Control Systems: A Framework for Analysis*. Cambridge: Harvard University Press; 1965.
- Simon HA. *The New Science of Management Decision*. United States: Harper and Row; 1960.
- Rockart JF. Chief executives define their own data needs. *Harv Bus Rev*. 1979;57(2):81-93.
- Alter S. A taxonomy of decision support systems. *Sloan Manage Rev*. 1977;19(1):39-56.
- Keen PGW. *Decision Support Systems: An Organizational Perspective*. Pearson: Addison-Wesley; 1978.
- Sprague RH. A framework for the development of decision support systems. *MIS Q*. 1980;4(4):1-26.
doi: 10.2307/248957
- McKinsey and Company. *Management Information Systems*. United States: McKinsey and Company; 1968.
- Shim JP, Warkentin M, Courtney JF, Power DJ, Sharda R, Carlsson C. Past, present, and future of decision support technology. *Decis Support Syst*. 2002;33(2):111-126.
doi: 10.1016/S0167-9236(01)00139-7
- Turban E, Aronson JE. *Decision Support Systems and Intelligent Systems*. 6th ed. United States: Prentice Hall; 2001.
- Howard RA, Matheson JE. Influence diagrams. In: Howard RA, Matheson JE, editors. *Readings on the Principles*

- and Applications of Decision Analysis*. Vol 2. California: Strategic Decisions Group; 1984. p. 719-762.
26. Zadeh LA. Fuzzy sets. *Inf Control*. 1965;8(3):338-353.
doi: 10.1016/S0019-9958(65)90241-X
27. Keeney RL, Raiffa H. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. United States: John Wiley and Sons; 1976.
28. Huele R. Neural Networks as Decision Support Systems New Tools for Handling Soil Contaminations. In: *Contaminated Soil '90*. Springer Netherlands; 1990:515-519.
doi: 10.1007/978-94-011-3270-1_117
29. Hertz DB. Risk analysis in capital investment. *Harv Bus Rev*. 1979;42(1):95-106.
30. Sittig DF, Wright A, Coiera E, *et al*. Current challenges in health information technology-related patient safety. *Health Informatics J*. 2020;26(1):181-189.
doi: 10.1177/1460458218814893
31. Kawamoto K, Houlihan CA, Balas EA, Lobach DF. Improving clinical practice using clinical decision support systems: A systematic review of trials to identify features critical to success. *BMJ*. 2005;330(7494):765.
doi: 10.1136/bmj.38398.500764.8F
32. Miller CC, Reardon MJ, Safi HJ. *Risk Stratification: A Practical Guide for Clinicians*. Cambridge: Cambridge University Press; 2001.
33. D'Amico AV, Whittington R, Malkowicz SB, *et al*. Biochemical outcome after radical prostatectomy, external beam radiation therapy, or interstitial radiation therapy for clinically localized prostate cancer. *JAMA*. 1998;280(11):969-974.
doi: 10.1001/jama.280.11.969
34. Arias E, Xu J, Kochanek KD. United States Life Tables, 2016. *Natl Vital Stat Rep*. 2019;68(4):1-66.
35. Johnson AEW, Pollard TJ, Shen L, *et al*. MIMIC-III, a freely accessible critical care database. *Sci Data*. 2016;3(1):160035.
doi: 10.1038/sdata.2016.35
36. Power DJ. *Decision Support Systems: Concepts and Resources for Managers*. United States: Greenwood Publishing Group; 2002.
37. Barnett GO, Cimino JJ, Hupp JA, Hoffer EP. DXplain. An evolving diagnostic decision-support program. *JAMA*. 1987;258(1):67-74.
doi: 10.1001/jama.258.1.67
38. Wallace A. *The Prodigy*. Hertford: Crossroad Press; 2017.
39. Weed LL. Problem-knowledge coupler. *MD Comput*. 1978;3(5):45-54.
40. Kaiser K, Miksch S. Modeling treatment processes using information extraction. In: Yoshida H, Jain A, editors. *Advanced Computational Intelligence Paradigms in Healthcare – 1 (Studies in Computational Intelligence)*. Vol. 48. Berlin: Springer; 2007:189-224.
doi: 10.1007/978-3-540-47527-9_8
41. Adadi A, Berrada M. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*. 2018;6:52138-52160.
doi: 10.1109/ACCESS.2018.2870052
42. Ganguly R, Singh D, Bose R. The next frontier of explainable artificial intelligence (XAI) in healthcare services: A study on PIMA diabetes dataset. *The Scientific Temper*. 2025;16(05):4165-4170.
doi: 10.58414/scientifictemper.2025.16.5.01
43. Turban E, Aronson JE, Liang TP. *Decision Support Systems and Intelligent Systems*. 6th ed. United States: Prentice Hall; 2007.
44. Garg AX, Adhikari NK, McDonald H, *et al*. Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: A systematic review. *JAMA*. 2005;293(10):1223-1238.
doi: 10.1001/jama.293.10.1223