Predictive AI Meets Evidence-Based Clinical Decision Support

Darena Solutions
MeldRx

Executive Summary

In the ever-evolving healthcare landscape, clinicians and organizations face the dual challenge of staying current with emerging clinical guidelines and integrating these updates seamlessly into patient care. Clinical Decision Support (CDS) tools have long served as a bridge, helping to operationalize evidence-based care within Electronic Health Records (EHRs). Now, the ability of Artificial Intelligence (AI) to analyze vast and complex datasets, both structured and unstructured, has expanded the possibilities for CDS. Predictive AI, a subset of AI, offers data-driven insights that promise to transform clinical workflows and decision-making processes. However, adopting Predictive AI for decision support comes with unique challenges, such as ensuring alignment with evidence-based care, integrating with existing workflows, earning clinicians' trust, and avoiding the pervasive issue of alert fatigue.

This white paper discusses the concept of Predictive AI and explores its potential within clinical decision support systems. It also examines key considerations for organizations looking to adopt this technology, providing an overview of essential aspects to evaluate. These include the regulatory landscape, the role of evidence-based care, the role of HL7® FHIR® for interoperability, and the concept of a multi-model AI solution tailored to healthcare. By understanding these interconnected areas, this paper aims to serve as an educational resource for healthcare organizations seeking to harness the transformative power of Predictive AI.

Terminology Clarification

Clinical Decision Support (CDS) vs. Decision Support Interventions (DSI)

In this white paper, we use the term **Clinical Decision Support (CDS)** to describe tools, systems, and processes that provide clinicians, patients, or other stakeholders with person-specific information and relevant knowledge intelligently filtered at appropriate times to enhance healthcare delivery. The new (b)(11) regulation in the US refers to such systems as Decision Support Interventions (DSI); our use of CDS aligns with its broader recognition in the healthcare and informatics communities. For the purposes of this document, CDS and DSI are conceptually equivalent, with an emphasis on supporting clinical decision-making.

Predictive Analytics and Predictive AI

Predictive **analytics** and **predictive AI** are sometimes used interchangeably in the industry. For the context of this white paper, **Predictive AI** refers to advanced technologies that support clinical decision-making by employing algorithms or models trained on historical data. These technologies extend beyond foundational predictive capabilities to include a range of functionalities, such as:

- **Prediction:** Estimating the likelihood of future outcomes based on current and historical data. For example, a hospital risk model estimates a 30% likelihood of 30-day readmission for a heart failure patient based on certain factors.
- **Classification:** Categorizing data into predefined groups. For instance, a deep-learning algorithm flags chest X-ray findings as "likely benign" or "likely malignant."
- **Recommendation:** Suggesting actionable steps or treatments. A decision support tool might suggest the optimal next step in therapy for a diabetic patient based on current A1C levels and medication history.

- **Evaluation:** Assessing patient data against clinical criteria or guidelines. For example, an automated system checks whether a patient meets the criteria for a new anticoagulant by comparing lab results to published guidelines.
- **Analysis:** Identifying patterns or trends that inform clinical insights. An AI system could detect a concerning trend in renal function (rising creatinine over three visits), indicating possible acute kidney injury.

Predictive AI "App" vs. Predictive AI "Solution"

The word "app" is often associated with a mobile or web-based software program in many contexts. However, when discussing Predictive AI in healthcare, "apps" can just as easily refer to behind-the-scenes services that deliver predictive insights—whether through APIs (e.g., CDS Hooks), embedded algorithms in electronic health record (EHR) systems, or more conventional apps. In other words, the distinction between a "Predictive AI app" and a "Predictive AI solution" can be subtle and is frequently blurred; these terms are often used interchangeably in the industry and in this white paper.

Evidence-Based Vs Predictive AI CDS

🦋 Integrating Predictive AI into Clinical Workflows

The successful adoption of predictive AI in clinical decision support depends on a core understanding of the clinical workflows today. At the heart of healthcare delivery are two key stakeholders: a **patient** and a **provider**. The provider's primary role is to make informed decisions based on the patient's inputs, such as their symptoms, medical history, and other relevant data. Patients, however, generate vast amounts of data, ranging from lab results and medications to device-generated data and questionnaires, that providers cannot realistically analyze manually within the constraints of clinical practice.

Decision Support Systems (DSS) have been integrated into electronic health records (EHRs) to bridge this gap. Traditionally, these systems rely on **evidence-based care**, which includes clinical guidelines and protocols developed through rigorous research. Evidence-based DSS provides **alerts or suggestions** to providers, supporting their decision-making processes. Providers can act on these suggestions, search for additional information, or dismiss them based on clinical judgment. These actions generate new data, which feeds back into the system, creating a continuous refinement loop. Furthermore, providers can offer feedback on the system's recommendations, enhancing its accuracy and usability over time.

For a **Predictive AI Clinical Decision Support (CDS)** system to succeed, it must seamlessly integrate into this established workflow without disruption. Predictive AI should enhance decision-making by offering dynamic, data-driven insights while aligning with the trusted framework of evidence-based systems.

It's also important to mention that Evidence-based care and predictive AI are not competing paradigms but complementary forces that together create a more robust clinical decision-making framework. For instance, evidence-based alerts might guide a clinician to screen a diabetic patient for retinopathy annually, while predictive AI models can layer on deeper insights, such as forecasting a risk of diabetic foot ulcers in six months based on wearable data trends and recent hospitalizations. By integrating predictive AI's nuanced, real-time analysis with the rigor and trust of evidence-based practices, clinicians gain a higher-resolution view of patient care, enabling proactive and personalized interventions.

Figure 1: Embedding Predictive AI into Clinical Decision Support (CDS)

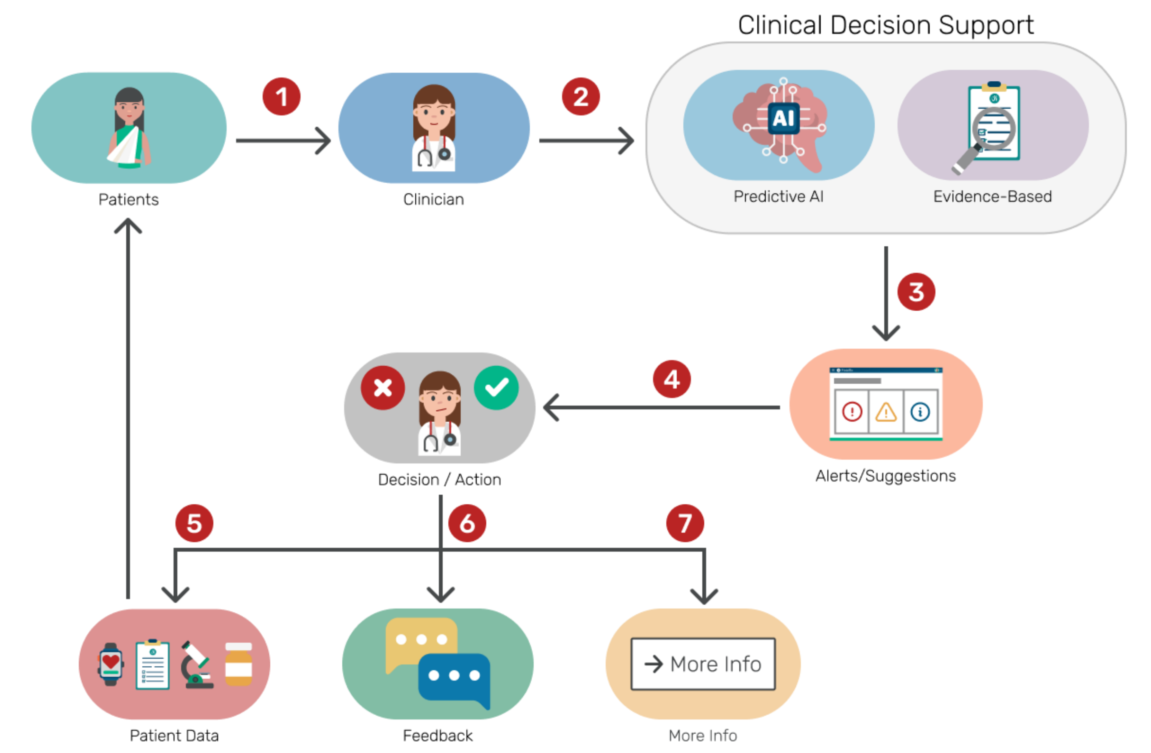

1. A patient visits their provider initiating a clinical workflow.
2. The clinician accesses the patient's chart, triggering the configured Clinical Decision Support (CDS) system based on **evidence-based** and **predictive AI** protocols.
3. The system generates alerts, insights or suggestions based on patient's data.
4. The clinician reviews the alerts and suggestions to make care decisions and take actions.
5. Actions (e.g. Tasks, Medications, Labs etc) generate new patient data.
6. Clinicians share feedback that helps improve protocols over time and reduce alert fatigue.
7. Clinicians may review more info such as clinical trials, research papers, source attributes of the CDS protocols before making their decision or taking action.

Understanding the nuances between these approaches is critical for aligning clinician decision-making with new AI capabilities. The following table illustrates the key distinctions, highlighting how each approach supports clinical decision-making in unique ways.

Aspect	Evidence-Based Care	Predictive AI
Source	Derived from clinical trials, observational studies, systematic reviews, and expert consensus.	Trained on real-world datasets (e.g., EHRs, CGM data) and learns patterns through machine learning algorithms.
Nature of Data	Based on rigorously collected, validated, and often static datasets.	Dynamically analyzes large, diverse, and often noisy datasets that may include biases or gaps.
Update Frequency	Periodic updates as new clinical evidence is published (e.g., annually).	Can have continuous updates as new data is ingested, though quality depends on access to sufficient and representative real-world data.
Personalization	Generalized; focuses on population-level effectiveness for the average patient.	Can personalize recommendations based on patient-specific data, though it struggles with edge cases or patients unlike those in its training set.
Validation	Rigorously validated through peer-reviewed research, randomized controlled trials, and meta-analyses.	Validation depends on model performance (e.g., accuracy, precision), external validation, and interpretability, which are not always robust.
Transparency	Highly transparent; the rationale for recommendations is clear and well-documented.	Often opaque ("black-box" models); explainability can be a challenge, and clinicians may not trust the rationale without clear evidence.
Scalability	Difficult to scale quickly for diverse populations or settings without extensive additional trials or expert reviews.	Easily scalable, but effectiveness depends on the quality and diversity of the training data.

Predictive AI Regulatory Landscape

Predictive AI, or AI in general, is not yet mandated for healthcare providers. Still, the industry is moving toward a landscape where its adoption may become an expected standard of care. Failure to integrate AI tools could be seen as a deviation from this standard, exposing providers to liability risks. A policy brief from Stanford's Institute for Human-Centered Artificial Intelligence highlights the benefits of traditional clinical decision support systems in reducing malpractice claims and preventing injuries. The brief cautions that not adopting emerging technological tools could eventually be considered a harmful decision. Similarly, the Coalition for Health AI (CHAI) emphasizes the need for harmonized standards and transparent reporting through its blueprint for trustworthy AI implementation. These developments underscore the growing expectation for healthcare providers to adopt AI tools to maintain and improve care standards.

Medicine is a science of exceptions, whereas predictive AI relies on identifying patterns in data. This divergence in approach is raising concerns about liability and trust when using predictive AI such as:

- **Incomplete Testing and Validation**: Predictive models may lack comprehensive clinical validation, creating uncertainty about their reliability in real-world settings.
- **Liability and Legal Implications**: If an AI-driven recommendation leads to an adverse outcome, questions arise about who bears the responsibility - clinicians, developers, or both.
- **Data Bias and Outliers**: Predictive AI models can systematically miss certain populations or fail to account for unusual cases, which remain a hallmark of medical practice.

Trust in predictive AI hinges on transparent regulatory frameworks and robust validation studies. Clinicians must be confident that the AI's recommendations are consistently reliable. Regulations surrounding Predictive AI in healthcare are rapidly evolving, reflecting the growing importance of transparency, accountability, and patient safety. Below is a list of key regulations shaping the U.S. market as of now:

1. HHS Final Rule (April 2024)

Overview: Under Section 1557 of the Affordable Care Act, the U.S. Department of Health and Human Services (HHS) issued a final rule to reinforce nondiscrimination protections, explicitly extending these protections to AI tools.

Scope of Nondiscrimination: The rule prohibits discrimination based on race, color, national origin, sex, age, or disability in health programs or activities.

Liability on Healthcare Organizations

- Covered entities (e.g., hospitals and clinics) are responsible for discriminatory outcomes arising from their use of AI.
- AI developers or vendors are not held liable for end-users' misuse or inappropriate implementation of AI tools.

2. ONC's Decision Support Interventions (DSI) Criterion [§170.315(b)(11)]

Purpose: Issued by the Office of the National Coordinator for Health IT (ONC/ASTP) to enhance transparency and trust in AI-driven decision support tools (Predictive DSIs) within EHR systems.

Key Requirements

- **Third-Party Integration:** EHR systems must allow users to integrate third-party clinical decision support (CDS) solutions, including predictive AI solutions, ensuring healthcare organizations can adopt and use external tools that meet their specific needs.

- **Source Attributes**: CDS Systems must reveal details on the origin, development, and intended use of predictive DSIs.

- **Risk Management**: EHR Vendors must implement and disclose practices addressing the validity, reliability, fairness, and safety of predictive DSIs.

EHR vendors were required to comply by December 31, 2024, even if they did not provide predictive DSIs themselves.

3. ISO/IEC 42001:2023 – Artificial Intelligence Management System

This international standard, while not specific to healthcare, guides organizations in establishing, implementing, maintaining, and continually improving an AI Management System (AIMS).

Core Focus Areas

- **Ethical Considerations**: Ensuring AI systems operate responsibly.
- **Transparency & Accountability**: Clear documentation of how AI systems function and who is responsible for them.
- **Risk Management**: Strategies for identifying and mitigating AI-related risks.

🦋 Implications for Independent AI Solution Developers

- While not directly regulated by the ONC's DSI criterion, developers will benefit from aligning with EHR vendors' compliance requirements to facilitate integration.
- Providing source attributes and risk management documentation will be crucial to meet healthcare organizations' expectations.
- Pursuing ISO/IEC 42001:2023 certification can underscore ethical and trustworthy AI development, making products more attractive to providers.

Predictive AI Solutions Marketplace

Thanks to the Cures Act and SMART on FHIR, the app ecosystem for EHRs has grown recently. However, EHR vendors still control access and often restrict key features such as write-back, limiting the development of advanced apps. With new AI capabilities emerging faster than any EHR can manage, the industry is finally shifting toward a "decoupling" trend. As independent developers build innovative AI apps, EHRs face more pressure to open their interfaces. This "decoupling" aligns with the ONC's (b)(11) regulation, which ensures healthcare organizations can choose any AI app without EHR vendor constraints. Essentially, (b)(11) acknowledges the need for innovation freedom so we can "mix and match" solutions.

In addition to the regulation, another trend is brewing in the AI solutions ecosystem. Data for training large language models has plateaued, so the next frontier is private and proprietary data, fueling more specialized AI applications. As large healthcare organizations recognize the strategic value of their data, they'll seek to build their own AI-powered solutions, creating even more demand for open standards and robust data exchange.

Ultimately, clinicians and researchers will source AI solutions from three primary avenues:

1. **EHR Vendors**: Many EHR systems already offer AI-powered decision-support tools, providing seamless integration within clinical workflows.

2. **Independent Developers**: Third-party AI applications deliver specialized functionalities to enhance capabilities beyond standard EHR offerings.

3. **In-House Development by Healthcare Organizations**: Leveraging proprietary data for customized AI, driving faster local innovation within organizations.

Regardless of the source, healthcare organizations must ensure that all AI applications comply with evolving regulations to mitigate risk and liability.

Role of HL7® FHIR® in Predictive AI

Data access is one of the most significant barriers to scaling AI in healthcare. Predictive AI applications rely on vast amounts of structured and unstructured data to train models. AI apps also require real-time access to patients' current data to deliver insights at the point of care. This introduces several challenges:

- **Limited Data Sharing:** Concerns about patient privacy, regulatory compliance, and data security make it difficult to integrate data from disparate systems for both training and real-time use.
- **Manual Data Entry:** Some AI applications depend on conversational interfaces or one-off data uploads to feed information. However, this approach is neither scalable nor feasible for real-time clinical decision-making.

The **HL7® FHIR® standard** has emerged as a transformative solution, addressing these challenges by enabling seamless data interoperability and integration. By leveraging FHIR, healthcare organizations and developers can create scalable, standards-based frameworks that eliminate many barriers to deploying predictive AI.

🦋 How FHIR Resolves Predictive AI Challenges

1. **Data Integration through FHIR**
    - **Standardized Data Models:** FHIR resources (e.g., Patient, Observation, Condition) ensure healthcare information is consistently represented and exchanged, simplifying data interoperability.
    - **APIs for Real-Time Exchange:** FHIR provides RESTful APIs that enable real-time data retrieval and updates, allowing predictive AI applications to process up-to-date clinical information.

- **Scalability Across Organizations:** With FHIR's widespread adoption and thorough documentation, developers can design predictive AI solutions that seamlessly adapt to multiple EHR systems and healthcare facilities.

By using FHIR, developers eliminate the need for custom integrations with individual hospitals or clinics. Instead, they rely on a shared, standards-based framework to scale AI solutions effectively.

2. Embedding Predictive AI into Clinical Workflows with SMART on FHIR

SMART on FHIR extends the capabilities of FHIR by embedding predictive AI into clinicians' existing workflows:

- **Contextual Launch:** SMART on FHIR allows AI applications to be launched directly within an EHR, passing patient context to eliminate manual entry.
- **Streamlined User Experience:** Integrated within the EHR interface, these apps reduce clinicians' cognitive burden and promote consistent adoption.
- **Minimized Workflow Disruptions:** SMART on FHIR ensures that accessing predictive insights feels intuitive, aligning seamlessly with evidence-based workflows.

This integration allows clinicians to view both patient data and AI-driven recommendations in a single system, significantly enhancing usability and efficiency.

3. Automating Clinical Decision Support with CDS Hooks

FHIR's potential is further expanded through **CDS Hooks**, which automates and embeds AI-driven insights directly into clinicians' workflows:

- **Real-Time Triggers:** CDS Hooks notify the EHR or external applications when clinical events occur, such as opening a patient chart or prescribing medication. This enables AI to deliver recommendations at the exact moment they're needed.

- **Seamless Embedding:** Predictive AI insights appear as “cards” or alerts within the clinician’s workflow, avoiding reliance on additional buttons or tabs.
- **Personalized Alerts:** Integrating predictive models into CDS Hooks allows recommendations to be tailored to individual patients, enhancing accuracy and relevance.

CDS Hooks delivers decision prompts exactly when and where they are needed, ensuring AI insights align with clinical workflows and support timely, evidence-based decisions.

- By addressing the challenges of data access, integration, and usability, **FHIR, SMART on FHIR, and CDS Hooks** collectively form a robust foundation for scaling predictive AI in healthcare. Together, they unlock the potential for real-time, personalized decision support while aligning with the clinician’s workflow, making predictive AI a practical and transformative tool in modern healthcare.

A Crash Course on AI Models

While ChatGPT and similar Large Language Models (LLMs) have taken center stage, Artificial Intelligence (AI) is far beyond the hype surrounding LLMs. AI models come in various shapes and sizes, each designed for different purposes. In healthcare, where data formats and clinical needs vary widely, different types of models play distinct roles. Choosing an AI solution is essentially choosing the model behind it. The model defines how well the solution aligns with clinical goals, whether ensuring accurate diagnoses, predicting patient outcomes, or aiding in decision-making. While a detailed discussion of all model types is beyond the scope of this paper, this section provides a quick crash course to understand the basics.

Three Broad Categories of AI Models

AI models can be grouped into **classical machine learning, Deep Learning, and Language Models.** Before diving into their specifics, it's important to understand two key concepts, Feature Engineering, and Training Paradigm, as they influence how these models process data and "learn." The concepts of Pre-Training and Fine Tuning are natural extensions.

Training Paradigm: This describes **how** a model is trained:

- **Task-Specific**: You build a model for a narrow goal (e.g., "predict sepsis risk").
- **End-to-End**: You feed raw data into a network, and it learns both feature extraction and prediction in one go (common in deep learning).

Feature Engineering: In simpler AI systems—particularly in classical machine learning—you select the specific data inputs (or "features") you believe are most relevant to predicting an outcome. For instance, to estimate the length of a patient's hospital stay, you might choose demographic data, diagnosis codes, and vitals. In deep learning, much of this manual "feature picking" is replaced by neural networks that detect the most important data signals on their own.

Pretraining and Fine-Tuning: This approach is typical for language models. You initially train a model on vast, general data (pretraining) and then adapt it to a specific task or domain (fine-tuning).

The table below compares the three broad AI Model categories across key aspects:

Aspect	Classical ML	Deep Learning	Language Models
Input Data Type	Structured/tabular	Unstructured (e.g., images, text)	Text (natural language)
Feature Engineering	Often manual	Automatically learned through neural nets	Automatically learned during pretraining
Training Paradigm	Task-specific	End-to-end	Pretraining + fine-tuning
Scale	Small	Medium to large	Large
Examples	Decision Trees, SVMs	CNNs, RNNs	GPT, BERT, PaLM

1. Classical Machine Learning

- **Feature Engineering**: Involves manually selecting which patient variables (e.g., lab results, vital signs) are relevant. You might decide that a patient's blood pressure, age, and cholesterol levels are features predictive of cardiovascular risk.
- **Training Paradigm**: Task-specific. A model is built for one objective at a time, such as classifying patient readmission risk.
- **Healthcare Example**: A simple logistic regression model to predict hospital readmissions might use carefully chosen features such as comorbidities, age, and discharge notes.

2. Deep Learning

- **Feature Engineering**: Instead of manually defining features, deep learning allows the neural network to learn relevant patterns directly from the data (e.g., from X-ray images or continuous EHR data).

- **Training Paradigm**: End-to-end. Data goes in, and the system "learns" to extract features and make predictions in one integrated process.
- **Healthcare Example**: A Convolutional Neural Network (CNN) diagnosing pneumonia from chest X-ray images. The model automatically learns what visual patterns indicate pneumonia, rather than relying on manually extracted features.

3. Language Models

- **Feature Engineering**: Typically learned automatically during a pretraining phase, especially for large-scale models.
- **Training Paradigm**: Pretraining on massive text corpora, then fine-tuning for specific healthcare tasks (e.g., summarizing clinical notes).
- **Healthcare Example**: A fine-tuned GPT-based system that helps physicians automatically draft patient discharge summaries or analyze clinical guidelines.

  In the context of language models, you may hear references to a model's "size," often measured by something called **parameters**. You can think of these parameters as the model's internal "knobs and dials" that adjust themselves based on the data the model sees. The more parameters a model has, the more nuanced its understanding of language, but may require more computer power. While LLMs are suitable for general purposes, a new concept of Small Language Models (SLM) is gaining traction, especially in healthcare.

🦋 Key Benefits of Small Language Models

- **Privacy:** Healthcare data demands the highest level of security. SLM can be deployed on-premises within tightly controlled networks, offering better data protection and easier compliance with stringent regulations.

- **Total Cost of Ownership:** LLMs require substantial computing resources, driving up costs and the carbon footprint. Smaller models can still deliver remarkable performance without expensive infrastructure.

- **Ease of Fine-Tuning:** Adapting a massive model for specialized healthcare tasks is often complex and requires advanced techniques like **Retrieval-Augmented Generation (RAG)**, where external knowledge sources are integrated to enhance the model's outputs, or function calls to refine its capabilities on an ongoing basis. In contrast, smaller LLMs can typically be thoroughly fine-tuned quickly and cost-effectively, precisely meeting an organization's needs. This advantage is especially crucial for healthcare applications, which demand granular domain expertise and real-time access to specific, reliable data.

🦋 Why Fine-Tuning Matters More Than You Think?

- **Full Customization:** Fine-tuning lets you integrate organization-specific terminology, protocols, and workflows into the model.

- **Improved Accuracy:** Domain-focused and even organization-focused training data often yield more reliable outcomes than generic "one-size-fits-all" approaches.

- **Reduced Reliance on External Workarounds:** While RAG and function-calling can enhance large models, smaller models can be refined at their core, leading to more consistent performance.

🦋 Key Takeaways for Healthcare

- **Different AI Tools for Different Tasks:** A simple classical ML model might be more transparent and easier to maintain for predicting patient admission rates. In contrast, a deep learning model could excel at analyzing X-ray images.

- **Data Quality and Structure Matter:** Classical ML often works well with tabular EHR data, whereas deep learning and language models thrive on unstructured data like clinical notes, audio recordings (e.g., doctor-patient consultations), or imaging studies.

- **Small vs. Large Language Models:** When dealing with text (e.g., patient histories, insurance forms), a small language model might suffice for specialized tasks, but large language models can handle a wider array of challenges if you have the computational resources.

Healthcare organizations can better identify the right Predictive AI solutions for their specific clinical or operational needs by appreciating these distinctions, especially around feature engineering and training paradigms.

A Multi-Model (Not Modal) Approach for Healthcare

The healthcare industry generates an enormous volume and variety of data—ranging from electronic health records (EHRs) and medical imaging to wearable device data, lab reports, and clinical notes. This diversity has created a demand for **multi-modal AI**, A multi-modal model is an AI model that processes and integrates multiple data types, such as text, images, and numerical values. A multi-modal approach is gaining traction for its ability to deliver holistic insights by leveraging the strengths of diverse data.

However, as the complexity of healthcare data increases, a more flexible and efficient paradigm emerges - **multi-model**. In this approach, multiple specialized models collaborate, each focusing on a specific task or data type. By dividing the workload and capitalizing on individual model strengths, multi-model AI delivers more comprehensive, accurate, and scalable solutions.

Much like healthcare relies on specialists, such as cardiologists, oncologists, and neurologists, working alongside primary care providers, multi-model AI avoids relying on a single "mega-model" to handle everything. Instead, specialized models excel in their respective domains, collectively offering richer and more reliable outcomes. For instance, a computer vision model can analyze medical images, a time-series model can process real-time vital signs, and a large language model (LLM) can summarize clinical notes or aid in patient communication. Together, these models can address a broader range of clinical challenges with precision and efficiency.

- Beyond accuracy, **multi-model AI** offers practical advantages. Smaller, specialized models require less computational power than large, general-purpose systems, reducing costs and improving scalability, which is especially important in resource-constrained healthcare settings. This combination of specialization, precision, and efficiency makes multi-model AI an essential paradigm for advancing healthcare solutions, aligning with the interdisciplinary nature of healthcare itself.

About Darena Solutions

For over 15 years, Darena Solutions has developed practical solutions to improve healthcare delivery, focusing on real-world challenges like improving data sharing between systems, streamlining clinical workflows, and ensuring compliance with regulatory standards.

MeldRx Overview

MeldRx is Darena Solutions' flagship offering, an FHIR-first backend-as-a-service platform designed to simplify the building and deployment of custom healthcare solutions for organizations of all sizes.

🦋 MeldRx Workspaces

The power of the MeldRx platform comes from its **workspaces**, which are FHIR-ready backend databases that can be provisioned in seconds and scaled on demand. These workspaces provide:

- **Full Data Control**: Organizations can upload and manage their data independently, enabling complete customization of solutions.

- **Apps Integration**: MeldRx offers a comprehensive developer portal, allowing organizations to easily configure their own apps or use apps published by third-party organizations with their workspaces.

Beyond its core functionality, MeldRx workspaces include other features that simplify data acquisition and expand their usability:

- **Linked Workspaces for Real-Time EHR Integration**

  MeldRx workspaces can be seamlessly "linked" to any EHR's FHIR servers, enabling real-time access to patient data through apps. Linked Workspaces' unique distributed architecture enables apps to support features such as managing custom data and overcoming EHR limitations (e.g., inability to write data back to the EHR).

- **Consumer-Mediated Exchange**

  MeldRx workspaces also facilitate easy integration of consumer-mediated data exchange, empowering organizations to simplify capturing patient data directly into the workspaces.

- **Regulatory Compliance**

  MeldRx is ONC-certified, providing healthcare organizations and EHR vendors the confidence to meet evolving regulatory requirements. Its adoption by multiple EHRs highlights its reliability and flexibility, making it a trusted regulatory compliance and innovation partner.

MeldRx CDS: Combining Predictive AI with Evidence-Based CDS

MeldRx Clinical Decision Support (CDS) is a **plug-and-play** component of MeldRx designed to streamline healthcare organizations' adoption of predictive AI in their current workflows. Through ongoing collaboration with EHRs, healthcare organizations, and AI solution developers, and guided by the key areas highlighted in this white paper, we continuously refine our approach. As we learn and adapt to real-world insights, we remain committed to ensuring practical and thoughtful adoption of AI in clinical workflows, balancing innovation with practicality.

MeldRx CDS leverages the core features of MeldRx workspaces and can be deployed as a standalone solution or integrated with any EHR. Additional key features include:

- **Unified Dashboard:** MeldRx CDS offers a fully configurable unified dashboard that consolidates alerts from evidence-based and AI-driven solutions. These solutions can be powered by third-party or in-house applications, allowing organizations to tailor the dashboard to their specific needs and workflows.
- **CDS Hooks**: MeldRx CDS uses CDS Hooks for alert delivery, actions, and collecting feedback. It can leverage CDS hooks even if the integrating EHR doesn't support them.
- **Real-Time EHR Integration**: MeldRx CDS can function as a standalone solution or seamlessly integrate with an EHR using SMART on FHIR.
- **Evidence-Based Alerts**: Built-in CQL designer allows the configuration of evidence-based rules and deploying them along with the predictive AI solutions
- **Multi-Model Enablement**: A robust developer portal for organizations to configure internal AI apps or use third-party solutions.
- **Feedback-Driven Alert Management**: A feedback system built to minimize alert fatigue.
- **Certified (b)(11) Compliance**: MeldRx was the nation's first certified (b)(11) solution.

Figure 2: Plug-and-Play Evidence-Based and Predictive AI CDS

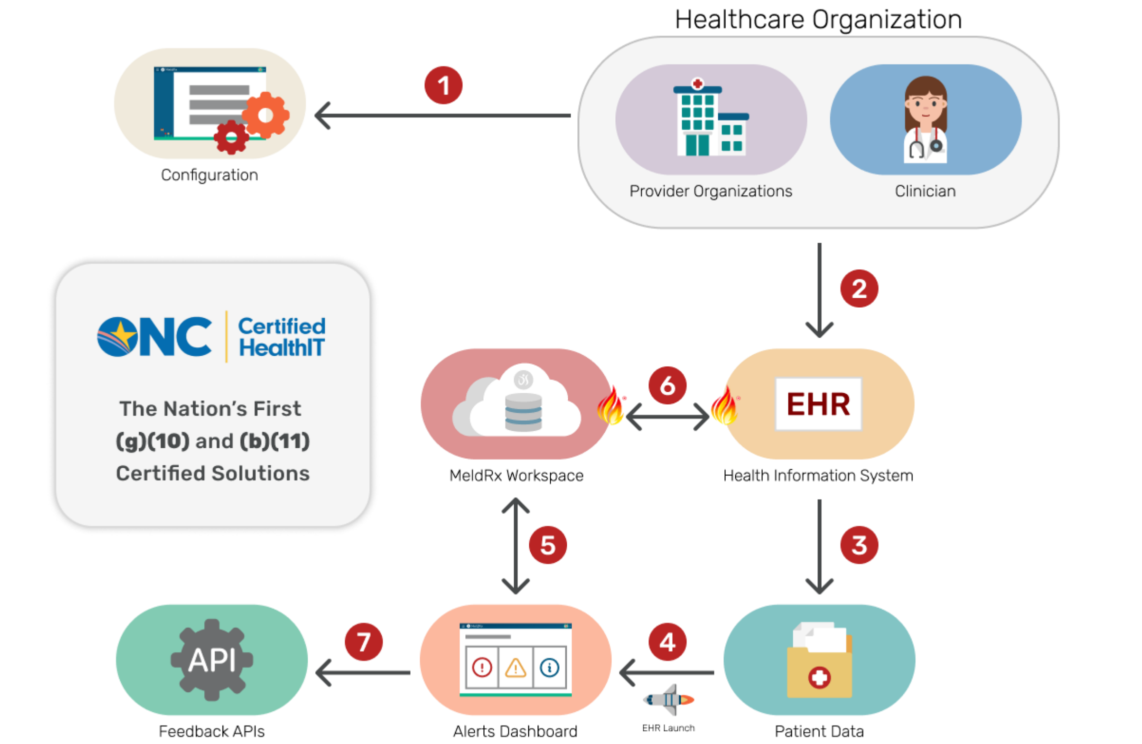

1. Organization configures the **evidence-based** or **predictive AI** CDS protocols in MeldRx.
2. A clinician logs in to their Health Information System (e.g. EHR).
3. The clinician opens a patient chart. The chart shows them a button to launch the alerts dashboard.
4. The clinician launches the dashboard, which displays alerts, insights, or suggestions based on configured protocols and the patient's data from the EHR and/or the MeldRx workspace.
5. The clinician reviews the alerts and suggestions to make care decisions and take actions. Data generated from actions is synced back to the MeldRx workspace.
6. Data generated from actions *may* also be synced back to the EHR.
7. Clinicians share feedback that helps improve protocols over time and reduce alert fatigue.

Call to Action: Evaluating Your Predictive AI Strategy

As healthcare organizations look to the future, taking a structured approach to AI adoption is vital. We invite you to:

- **Review Your Readiness**
  Assess whether your infrastructure, data governance, and clinical workflows are prepared to support predictive AI at scale.

- **Engage Stakeholders**
  Involve clinicians, IT teams, and leadership in identifying high-impact opportunities where predictive AI could improve outcomes and reduce costs.

- **Leverage a Trusted Partner**
  Consider partnering with Darena Solutions. Our MeldRx CDS platform offers a plug-and-play solution for unifying AI-driven and evidence-based insights in real-time integrated with your EHR or other systems.

Following the framework outlined above and leveraging tools like MeldRx, you can set a clear path for successful AI adoption. Our team is here to support you in defining a roadmap, rolling out targeted pilots, and scaling solutions that combine clinical evidence with advanced predictive analytics. Reach out today at [www. meldrx.com](www.meldrx.com) to learn how MeldRx can help you realize your AI aspirations and transform healthcare delivery.