Clareus Scientific Science and Engineering
Volume 3 Issue 3 May 2026
DOI: 10.70012/CSSE.03.078
ISSN: 3065-1182

Artificial Intelligence to Help Reduce Delays and Bottlenecks in the Care Pathways for Rare Diseases: An Integrated Practitioner–Patient Approach

Citation: Robin Vivian. "Artificial Intelligence to Help Reduce Delays and Bottlenecks in the Care Pathways for Rare Diseases: An Integrated Practitioner–Patient Approach". Clareus Scientific Science and Engineering 3.3 (2026): 15-27.

Article Type: Review Article

Received: May 15, 2026

Published: May 29, 2026

Robin Vivian*

Perseus Laboratory, University of Lorraine, Metz, France

***Corresponding Author:** Robin Vivian, Perseus Laboratory, University of Lorraine, Metz, France.

Abstract

Rare diseases affect millions of people worldwide. They are characterised by prolonged diagnostic uncertainty, fragmented care pathways and significant logistical and informational barriers. Current hospital information systems and conventional artificial intelligence (AI) approaches are primarily designed for common conditions and remain poorly adapted to the sparse, heterogeneous and longitudinal nature of rare disease data. This article critically reviews recent advances in AI applied to rare diseases, including electronic health record (EHR) analysis, anomaly detection, multimodal phenotyping, large language models (LLMs) and agent-based systems. Based on these developments, the paper presents a practitioner–patient framework centred on the continuous identification of unusual data, rather than direct automated diagnosis. This framework combines anomaly detection models, natural language processing, phenotypic matching based on the Human Phenotype Ontology (HPO), multimodal data integration and explainable AI mechanisms, supporting clinical vigilance throughout the care pathway. In parallel, a secure, multimodal AI assistant is introduced to empower patients by simplifying medical information, improving coordination and reducing administrative burdens. Unlike existing approaches, which focus primarily on early diagnosis, the proposed framework aims to minimise bottlenecks throughout the patient journey, while ensuring human clinical oversight is maintained and compliance with emerging regulatory requirements, such as the GDPR and the EU AI Act, is achieved. The article discusses the technical, ethical, organisational and regulatory challenges associated with such systems, and highlights the potential of AI to improve the timeliness of diagnoses, care coordination and patient autonomy in the context of rare diseases.

Keywords: rare diseases; artificial intelligence; electronic health records; anomaly detection; multimodal phenotyping; large language models; explainable AI; patient pathway optimisation; clinical decision support; healthcare interoperability

Introduction

Rare diseases affect around 300 million people worldwide and over 3 million people in France. They are characterised by extreme clinical diversity and low prevalence. They are often characterised by prolonged diagnostic uncertainty (averaging 5 to 10 years), as well as logistical and informational barriers, such as multiple consultations, redundant tests, fragmented medical records and a lack of understanding of treatment options among patients and their families. These challenges are exacerbated by the rarity of these conditions and highlight the urgent need for tools that can detect atypical signs early and streamline care pathways. While considerable efforts are being made in the medical sector, current information systems, which are designed for common conditions, struggle to exploit longitudinal data fully. This data is often heterogeneous, incomplete and inaccurate. Artificial intelligence (AI) is emerging as a complementary tool to enhance practitioners' capabilities, such as the detection of hidden patterns and the analysis of multimodal data, and to support patients in understanding diagnoses and care planning.

This article offers a critical review of AI applications in rare diseases, focusing on recent advances in electronic health record (EHR) analysis, multimodal phenotyping and agent-based systems. We then explore how an integrated practitioner–patient–AI approach centred on the continuous detection of atypical data could minimise bottlenecks throughout the care pathway without ever replacing medical expertise.

State of the art

Due to their low prevalence (less than 1 in 2,000) and clinical heterogeneity, rare diseases pose major computational challenges for traditional AI systems. Despite the increasing digitisation of health data, current tools, which are designed for common conditions, struggle to utilise electronic health records (EHRs) effectively. These records contain a variety of information, including structured data such as diagnostic codes and laboratory results, semi-structured data such as free-text clinical notes, and unstructured data such as images and genomic data. This information is often fragmented across institutions and incomplete. This complexity creates an extreme imbalance between classes (few positive cases versus a large amount of negative or unreliable data), making traditional machine learning approaches unsuitable without specific adaptation. This uncertainty leads to an increase in specialist consultations and complex inter-specialist coordination, placing a significant cognitive, emotional, and administrative burden on patients, their families, and practitioners (Visibelli et al., 2023; Groza et al., 2026).

The digitisation of information was intended to solve many of the problems experienced by patients during their care journey. Information technology has indeed partially fulfilled this role in terms of digitising examinations, archiving data and providing access to information. However, although processing tools handle conventional data correctly, problems remain in detecting weak and atypical signals within large amounts of heterogeneous, often incomplete and noisy longitudinal data. Electronic medical records (EMRs) contain three types of information: structured (such as test results), semi-structured (free-form notes) and unstructured (multimodal data such as images or phenotypes). The rarity of positive cases creates an extreme imbalance between classes, rendering classic machine learning approaches unsuitable without modification. Furthermore, a lack of interoperability between hospital systems, confidentiality constraints (e.g. GDPR[1], HIPAA[2]) and the need for high explainability for clinical adoption constitute major technical obstacles (Gebeyehu et al., 2026). Despite these challenges, artificial intelligence has made significant advances in recent years, primarily in the areas of early diagnosis and image analysis. This section reviews these contributions across four main areas: 1) EHR[3] analysis and anomaly detection; 2) multimodal phenotyping (i.e. measuring visible characteristics); 3) large language models (LLMs) and agent-based systems; and 4) persistent limitations and prospects for supporting the entire care pathway.

The IT context of barriers to the care pathway in rare diseases

Contemporary hospital information systems, the legacy of decades of substantial investment in the most common conditions, are firmly grounded in a logic of volume and standardisation. This approach often struggles to address the clinical, biological, psychosocial and organisational complexities of rare diseases, which are frequently overlooked in terms of research, treatments and resources. Within these systems, data relating to medicine, administration, biology, genetics, imaging, follow-up and quality of life is scattered

across multiple public and private institutions and siloed medical specialities. This data is also found in a variety of formats, including unstructured free text, ICD-10/ICD-11 codes[4], which are sometimes poorly documented, digital imaging files (DICOM[5]), raw or interpreted genomic sequencing results, molecular biology reports, outpatient consultation forms, anaesthesia records, surgical reports, medication prescriptions and nursing care protocols. This makes creating a longitudinal, coherent and reliable patient pathway for diagnostic, therapeutic monitoring or clinical research purposes a technological, organisational and human challenge.

From an algorithmic perspective, there are several obstacles that largely explain why current AI systems struggle to transition from one-off detection to providing continuous and reliable support throughout the care pathway for rare diseases.

a. ***The scarcity of positive training examples, which limits supervised approaches***. Rare diseases are characterised by an extremely low prevalence (often less than 1 in 2,000 people), which results in extreme class imbalance in training datasets (Visibelli et al., 2023). Conventional supervised classification algorithms, such as deep neural networks, random forests and Support Vector Machines (SVMs), typically require hundreds or thousands of positive examples to converge towards robust generalisation. However, in this field, the number of confirmed cases is often limited to a few dozen per hospital. This scarcity can lead to severe overfitting and low sensitivity to ultra-rare cases. Various compensatory strategies have been explored, including few-shot learning, meta-learning and federated learning (Zhu et al., 2026; Kolekar et al., 2025). However, these strategies are still limited by fragmented data across institutions and the absence of large-scale, harmonised cohorts. In the context of EHRs, this results in recall performance often being below 30% for the rarest conditions, making the systems unreliable for the early detection of rare diseases (Salmi et al., 2024).
b. ***The temporal variability of symptoms (gradual onset, intermittent episodes)***. Unlike stable chronic conditions, rare diseases often follow non-stationary trajectories, with symptoms appearing gradually over several years and experiencing intermittent exacerbations or spontaneous remissions. Standard models such as Short-Term and Long-Term Memory, Gated Recurrent Units and classical transformers frequently assume temporal regularity or uniform sampling. However, this does not correspond to the reality of irregular medical visits and variable time intervals in EHRs (AlSaad et al., 2024; Sreenivasan et al., 2026). Recent temporal approaches (e.g. temporal self-attention, temporal graphs and longitudinal transformers) attempt to model these irregularities via temporal vector representations or adapted attention masks. However, they remain sensitive to misalignments between clinical events and information loss over long time horizons (Moglia et al., 2025). In rare diseases, this temporal variability makes it particularly challenging to identify hidden causal patterns, as atypical signals may be scattered over several years and obscured by frequent intercurrent events.
c. ***Noise and missing data***. Electronic health records (EHRs) for patients with rare diseases may contain a high proportion of missing data (up to 40–60% for certain longitudinal fields) and errors (e.g. coding errors, ambiguous clinical notes and incomplete test results). These issues are exacerbated by patients visiting multiple healthcare facilities and the resulting fragmentation of records (Patharkar et al., 2024). While conventional imputation techniques (e.g. averaging, K-nearest neighbours (KNN), and forward-fill) and generative methods (e.g. generative adversarial networks (GANs) and variational autoencoders) can improve data completeness, they often introduce artefacts that can bias subsequent predictions, particularly in contexts of low prevalence (Digitale et al., 2025). Textual noise (unstructured notes) also necessitates robust natural language processing (NLP) pre-processing, which increases computational complexity and the risk of error propagation in multimodal pipelines. Handwriting recognition systems are not even worth mentioning here; OCRs certainly remain unable to decipher doctors' handwriting.
d. ***The need for multimodal integration (text + image + genomics)***. Data relevant to rare diseases are inherently multimodal, comprising textual clinical notes (handwritten or typed), medical images (e.g. MRI, X-ray, bone scintigraphy and facial phenotyping), genomic data (e.g. VCF variants) and longitudinal measurements (e.g. biomarkers). Integrating these heterogeneous modalities presents challenges in terms of temporal alignment, dimensionality, and handling missing data. Although early or late fusion approaches using multimodal transformers or cross-attention mechanisms demonstrate improved performance, they are more sensitive to incomplete modalities (Viswan et al., 2026; Doan et al., 2026). In the context of rare diseases, where a single modality (e.g. imaging) may be inadequate, the absence of robust fusion can significantly diminish the system's overall sensitivity

and complicate the interpretability of its decisions.

e. ***The requirement for traceability and explainability to comply with regulatory standards (the EU AI Act, transparency requirements)***. AI systems applied to patient pathways are classified as 'high-risk' under the European Union's AI Regulation (the EU AI Act, due to be implemented from 2026), which requires full traceability of decisions and technical and human explainability, as well as mandatory human oversight. 'Black-box' models (e.g. deep learning and LLM agents) struggle to provide explanations that are both locally accurate and globally consistent (e.g. SHAP, LIME, attention maps or agentic reasoning chains). However, without explainability, clinical integration remains limited and the risk of non-compliance with regulations is high (Krejcar et al., 2026; Seifi et al., 2025). In the context of rare diseases, where ethical and medical issues are particularly sensitive, a lack of traceability hinders post-hoc auditing and undermines the trust of practitioners and patients.

These five algorithmic barriers are not independent; they interact and amplify each other's effects, making the holistic design of future AI architectures essential (Groza et al., 2026). This paper therefore proposes a hybrid practitioner–patient framework centred on the continuous detection of atypical data, with the aim of overcoming these constraints. The authors further propose a conceptual framework centred on the patient journey. They advocate organising AI around a 'patient–clinician–AI' triad, in which AI would intervene at every stage of the patient journey, from early suspicion via screening of electronic health records and support for differential diagnosis to optimisation of care coordination, enrolment in clinical trials and personalised longitudinal follow-up. Unlike other models, this one explicitly addresses the removal of logistical and informational bottlenecks throughout the patient journey, not just in the diagnostic phase.

EHR analysis and anomaly detection approaches

Lin et al. (2023) developed a system that combines natural language processing (NLP) and phenotypic matching (the Symptoma approach). This system has been successfully applied to anonymised retrospective electronic health records (EHRs). This combination is one of the cornerstones of early detection systems for rare diseases because it transforms unstructured textual data into standardised phenotypic profiles that can be processed by similarity algorithms.

The process consists of two closely linked main stages:

- ***First stage: phenotypic extraction via NLP***. Natural language processing (NLP) analyses unstructured clinical notes, such as free-text reports, medical histories and observations, as well as semi-structured data from electronic health records (EHRs). It uses Named Entity Recognition (NER) techniques to identify relevant medical terms such as symptoms, clinical signs and differential diagnoses. BERT-type models or specialised pipelines (such as spaCy/medspaCy or UMLS-based approaches) detect not only entities, but also their context, including negation ('absence of', 'without'), assertion ('suspected', 'probable') and experiencer ('does the symptom concern the patient or a family member?'). This 'deep phenotyping' phase extracts hundreds of potential mentions while filtering out irrelevant information and false positives.
- ***Step two: normalisation and phenotypic matching***. The terms extracted by NLP are normalised and aligned with a standardised ontology, primarily the Human Phenotype Ontology (HPO). This mapping process converts free-form expressions such as 'severe chronic fatigue' and 'proximal muscle weakness' into specific, interoperable HPO codes (e.g. HP:0003326 for 'proximal limb muscle weakness'). Phenotypic matching involves comparing the patient's HPO profile (a vector of weighted terms) with profiles of rare diseases stored in knowledge bases such as Orphanet, OMIM and the Mondo Disease Ontology.

The algorithm calculates phenotypic similarity using either semantic metrics, such as Lin similarity or cosine similarity on phenotypic embeddings, or supervised ranking models. The result is a probability score for each target rare disease. In a recent study (Lin et al., 2023), this approach facilitated the screening of 350,116 electronic health records (EHRs) spanning 15 years in under a month, identifying 104 patients suspected of having Pompe disease (type II glycogen storage disease) with remarkable efficiency. On average, doctors needed to review just 5.47 records to identify a probable case. This combination of NLP and phenotypic matching offers three key advantages: (1) it exploits the wealth of textual data that is often overlooked in structured EHRs alone; (2) it ensures interopera-

bility through the HPO ontology; and (3) it addresses class imbalance by focusing on similarity rather than traditional binary classification. However, the method remains sensitive to the quality of clinical notes and the completeness of the HPO ontology for ultra-rare conditions. Recent studies (Faviez et al., 2025; Mao et al., 2025; Germain, 2025) confirm that enhancing the NLP phase (via enriched UMLS+ models or LLMs) can greatly improve the sensitivity of phenotypic matching. For example, it can increase from 49% to 95% for certain rare syndromes while maintaining high specificity.

In summary, the combination of NLP (for contextual extraction) and phenotypic matching (for standardisation and semantic comparison) turns EHRs into a resource that can be used for early detection. This paves the way for large-scale screening systems that do not require large amounts of labelled training data.

More recent studies have focused on anomaly detection (or rare event detection) in rare medical contexts. In a comprehensive review published in PLOS ONE, Gebeyehu et al. (2026) analyse current machine learning practices for detecting rare medical events in EHRs. They emphasise the effectiveness of unsupervised or weakly supervised methods such as variational autoencoders (VAEs), isolation forests, subspace methods and temporal graph-based approaches. These techniques can identify atypical trajectories without requiring a large number of positive examples by exploiting natural class imbalances. Systems such as InfEHR, for example, or pipelines based on temporal knowledge graphs, link clinical events scattered across time and between specialities, revealing hidden causal patterns. Recent work by Boardman-Pretty et al. (2026) evaluates case-finding algorithms such as MendelScan on structured EHRs for 34 rare diseases, demonstrating the ability to prioritise patients for in-depth clinical review.

From a technical perspective, these approaches typically use event embedding models, temporal transformers, or knowledge graphs that incorporate the Human Phenotype Ontology (HPO). Challenges in this area include scalability across large volumes of hospital data and managing noise in clinical notes. Nishat (2025) emphasises the importance of NLP for extracting phenotypes from unstructured notes, offering practical applications in the diagnosis of diseases such as transthyretin amyloid polyneuropathy (vATTR) and Pompe disease.

Multimodal phenotyping and specialised tools

Facial and multimodal phenotyping is a well-established field of AI in the context of rare diseases. Tools such as Face2Gene and GestaltMatcher use convolutional neural networks (CNNs) and structured phenotypic databases to analyse facial images and assist in the diagnosis of genetic syndromes. Real-world evaluations (2024–2025) report top-three accuracy rates of between 58% and 77%, with better performance for rare syndromes than ultra-rare ones. These systems integrate with clinical workflows via APIs and combine computer vision with textual data from electronic health records (EHRs). Multimodal approaches go beyond the face to fuse medical images, genomic data, clinical descriptions, and in some cases, data from wearable devices (such as smartwatches), via late fusion models or multimodal transformers. These architectures provide a unified representation of the patient, improving sensitivity to subtle patterns such as discrete morphological anomalies and genotype-phenotype correlations. Deep learning-based pipelines applied to imaging (X-rays and MRIs) can help to detect skeletal or visceral abnormalities in syndromes such as Jeune syndrome (asphyxiating thoracic dysplasia). Integration with ontologies (HPO and the Mondo Disease Ontology) facilitates semantic reasoning and interoperability.

LLMs and agent-based systems: towards augmented and explainable reasoning

The emergence of large language models (LLMs) and agent-based architectures is a significant advancement in the field of medical informatics for rare diseases. In a study published in JAMA Network Open, Shyr et al. (2025) evaluated LLMs on complex cases from the Undiagnosed Diseases Network (UDN). ChatGPT achieved a diagnostic rate of 13.3%, compared to 5.6% for historical clinical reviews, while providing useful differential diagnoses in 23.3% of cases. Further research indicates that integrating LLMs with bioinformatics tools, such as Exomiser, can increase combined accuracy to approximately 30%. Zhao et al. (2026) proposed DeepRare, a multi-tool agentic system published in Nature. Across 163 real-world clinical cases, DeepRare achieved an accuracy rate of 64% for the top 1 result (compared to 55% for senior clinicians using search tools) and 78.5% for the top 5 results. Its architecture is based

on a central LLM host (DeepSeek-V3 by default), which is surrounded by over 40 specialised agents, including a symptom extractor, a phenotype analyser based on HPO, a genotype analyser of VCF variants, a knowledge base search agent for PubMed and Orphanet, a self-reflection agent, and a reasoning chain generator. DeepRare can process heterogeneous inputs, such as free text, structured HPO terms and VCF files, and produce ranked diagnostic hypotheses with traceable reasoning chains that are linked to verifiable sources. Explainability has been validated at 95.4% by clinical experts, highlighting the value of agent-based approaches for traceability — an essential criterion in a regulated context. The three-tier architecture (LLM host, specialised agents and external knowledge sources) enables dynamic knowledge updates.

Current limitations and unresolved technical challenges

Despite this progress, several technical and methodological limitations remain.

Firstly, the majority of systems remain focused on the initial diagnosis and under-utilise the longitudinal and patient-centred dimensions, such as support with administrative procedures, optimised scheduling of tests and providing tailored explanations of reports in plain language.

Secondly, biases linked to the training data (e.g. underrepresentation of diverse ethnic populations or regions with low digitalisation) and a lack of robustness in the face of incomplete or inaccurate data pose significant risks (Gebeyehu et al., 2026; Nishat, 2025).

Thirdly, integrating systems into existing clinical software poses challenges in terms of interoperability (FHIR and HL7 standards), confidentiality, and acceptance by practitioners. Agent-based systems are promising but raise issues of computational cost, latency, and maintaining specialised agents.

Finally, prospective multicentre clinical evaluation remains rare, with most studies being retrospective or conducted on selected cases. The lack of unified evaluation standards for agent-based systems hinders the transition to clinical scale, as there is no quantitative measurement of explainability, actual impact on diagnosis time, reduction in unnecessary consultations or cost-effectiveness.

In summary, substantial progress has been made in the early detection of health issues through the analysis of electronic health records (EHRs), multimodal phenotyping and agent-based reasoning, all of which are augmented by large language models (LLMs). However, the transition to AI that is fully integrated into the entire care pathway — capable of continuously detecting atypical data, alerting clinicians without replacing their judgement and providing patients with practical support throughout their journey — remains largely unexplored in the field of medical informatics. Against this backdrop, this paper proposes a hybrid clinician–patient framework centred on minimising bottlenecks.

Continuous detection of atypical data and a non-diagnostic clinical alert system

General principle: vigilance support centred on atypicality

The detection of atypical data forms the central methodological basis of the proposals presented in this article. Unlike traditional diagnostic support approaches, which aim to identify a specific pathology directly, this approach focuses on highlighting weak, rare or inconsistent signals within the patient's longitudinal data without ever providing a diagnosis. The aim is to support clinical vigilance by drawing the practitioner's attention to atypical care pathways, which may suggest a rare disease or a blockage in the care pathway, while respecting the clinician's role in decision-making. Here, atypicality is defined in a multidimensional way: it may correspond to an unusual combination of symptoms; biological progression inconsistent with established diagnoses; an accumulation of subtle signs spread over time; or inconsistency between different data modalities (e.g. textual, biological, genomic or imaging). This approach is particularly well-suited to rare diseases, which are characterised by the temporal dispersion of signals, extreme class imbalance, and fragmented medical records.

Technical architecture for the detection of atypical findings

From an algorithmic perspective, the detection process is based on a hybrid architecture that combines several complementary components.

The first anomaly detection module analyses structured and semi-structured longitudinal data from electronic health records, such as laboratory results, ICD codes, prescriptions and the chronology of procedures. This module employs unsupervised or weakly supervised methods, such as variational autoencoders and temporal graphs. These methods are capable of modelling the distribution of typical clinical trajectories and identifying significant deviations, even when large volumes of labelled positive cases are unavailable. In parallel, a natural language processing (NLP) module extracts relevant information from unstructured clinical notes and standardises it using the Human Phenotype Ontology (HPO). This phenotypic matching enables the patient to be represented by a standardised phenotypic profile that is interoperable with rare disease knowledge bases such as Orphanet, OMIM and the Mondo Disease Ontology. The signals detected by the statistical and phenotypic modules are enriched by patient-reported data (e.g. patient-reported questionnaires) and, where available, by data from connected devices.

All this information is then aggregated to produce a multidimensional atypicality score that incorporates the statistical rarity, temporal consistency, and multimodal consistency of the observations. When this score exceeds an adaptive threshold, an alert is generated without any diagnostic hypotheses being formulated.

Clinical alert system integrated into the care pathway

The detection of atypical findings is implemented as a non-intrusive clinical alert system that is integrated into existing clinical software via interoperability standards, such as FHIR[6]. Operating in the background, the system queries medical record data in real time or at regular intervals depending on the capabilities of the hospital information systems. When an unusual pattern is identified — such as a combination of signs consistent with a rare disease that had not previously been considered, atypical biological progression or persistent discrepancies between symptoms and diagnoses — a discreet alert is displayed on the clinician's interface. This takes the form of a visual indicator (an icon or a side panel), which signals the presence of atypical data and prompts a clinical review. The alert is accompanied by a concise, traceable explanation generated by an explanatory module based on a large language model. This summarises the elements that led to the detection without revealing the model's internal details.

The system does not prescribe any action, nor does it prioritise diagnostic hypotheses. It merely signals a deviation from usual clinical trajectories, leaving the clinician with full responsibility for interpretation and deciding on the course of action. Each alert can be validated, ignored, or rejected by the clinician, who can provide feedback that can be utilised in a federated local learning framework without sharing sensitive data.

Traceability, explainability and regulatory compliance

The entire process is meticulously designed to comply fully with the elevated standards of traceability, explainability and human oversight that major European regulatory frameworks strictly require. These frameworks include the General Data Protection Regulation (GDPR), which safeguards personal data privacy, and the Artificial Intelligence Act (EU AI Act), which classifies medical AI systems used in patient care as high risk and imposes strict requirements for transparency, accountability and continuous human supervision. These regulations emphasise the need for robust mechanisms to ensure that AI systems remain transparent, accountable and under human control at every stage — especially when dealing with vulnerable populations affected by rare diseases, where missteps could have profound ethical, medical and legal consequences.

Every alert generated by the system undergoes a rigorous auditing process. Each alert is automatically time-stamped with high-precision metadata capturing the exact timing of its generation. It is also explicitly linked to every relevant data element that triggered it, ranging from structured entries such as laboratory results and ICD codes to semi-structured phenotypic profiles derived from the

natural language processing of clinical notes, temporal graphs indicating atypical trajectories or multimodal signals from imaging or patient-reported outcomes. This detailed linkage supports comprehensive traceability and allows for effective post-hoc verification, error analysis and continuous system optimisation, while fully respecting data confidentiality requirements and enabling independent regulatory audits whenever necessary. Moreover, the reasoning chains that accompany each alert, generated via a dedicated explanatory large language model module, are carefully simplified and rendered highly accessible and comprehensible to end-users in clinical practice. These chains are customized to the practitioner's workflow and specialty, using language that is both precise and free of unnecessary complexity, while incorporating visual aids or structured summaries where appropriate to enhance understanding. By doing so, the system actively promotes a culture of trust and confidence among medical professionals, which is essential for encouraging the adoption and routine use of such innovative tools in real-world hospital environments and for fostering long-term acceptance by the clinical community.

Moreover, the reasoning chains that accompany each alert are generated via a dedicated explanatory large language model module. These chains are carefully simplified to make them highly accessible and comprehensible to end users in clinical practice. Customised for the practitioner's workflow and specialty, these chains use precise, uncomplicated language and incorporate visual aids or structured summaries where appropriate to enhance understanding. This approach actively promotes a culture of trust and confidence among medical professionals, which is essential for the routine adoption of such innovative tools in real-world hospital environments and long-term acceptance by the clinical community. In line with this, the alert system focuses specifically on providing vigilance support and identifying atypical data patterns, rather than engaging in automated decision-making or providing prescriptive recommendations. This distinction is fundamental, as it ensures that the AI is strictly an augmentation to clinical expertise rather than a replacement. This aligns perfectly with the ethical principles that underpin the entire proposed practitioner–patient integrated approach, ensuring that human judgement remains the ultimate authority at every step of the care pathway.

As a result, this alert system, which is built on the innovative principle of continuously detecting atypical data throughout the patient's longitudinal record, serves as a central and foundational pillar within the broader framework presented in this article. It enables the seamless integration of advanced data analysis capabilities into everyday clinical operations.

Ultimately, it demonstrates profound respect for the indispensable role of medical expertise, while also contributing to the streamlining and optimisation of care pathways. This is particularly beneficial for patients with rare diseases or those experiencing prolonged diagnostic uncertainty, as it reduces delays, enhances coordination and improves outcomes without compromising the human element of healthcare delivery or undermining the central position of practitioners and patients in decision-making.

Multimodal patient-centric AI assistant

This project aims to develop a secure, chatbot-based, multimodal AI assistant that can be accessed via a mobile or web application. Its hybrid architecture combines a large language model specialising in medical vocabulary with speech recognition and image analysis capabilities, enabling it to process scanned or photographed medical documents. The assistant will support people with rare diseases throughout their care journey, reducing informational, logistical and administrative barriers. Compliant with the GDPR and the AI Act, it operates with explicit consent and guarantees end-to-end data encryption, while avoiding the permanent storage of sensitive information on servers.

Key features

- ***Explanation of medical documents***: The assistant rephrases test reports, genetic results and consultation letters, presenting them in simple, accessible language tailored to the patient's level of health literacy. They also generate explanatory diagrams, infographics or animations to clarify complex concepts such as anatomy, disease mechanisms and treatment impacts.
- ***Natural language interaction***: Patients can ask questions using voice commands, such as "What does this mutation mean?" or "Why is this test important?", and the assistant will respond in a contextualised and reassuring manner while highlighting the potential risks of misinterpretation.

- ***Logistical optimisation***: The assistant uses a constraint-based scheduling algorithm to organise appointments and tests. This takes into account medical constraints (e.g. deadlines and prerequisites), centre availability, the patient's geographical location, and their preferences. It sends automatic reminders via notifications or voice messages, monitors compliance through questionnaires or by integrating data from connected devices (smartwatches, etc.) and offers alternative solutions in the event of delays or missed appointments. It also informs the attending doctor, with the patient's consent.
- ***Smart referral***: The assistant uses an up-to-date database of rare disease centres of expertise (Orphanet, Filières de Santé) to suggest geolocated referrals tailored to the suspected or diagnosed condition. Inspired by the French AIDY project for phenotyping, it goes further by incorporating personalised, context-specific recommendations.
- ***Integrated and secure system***: With shared and revocable consent, the platform links the alert system for practitioners and the assistant for patients. It relies on a secure, GDPR-compliant database and standardised FHIR exchanges, as well as an agent-based orchestration engine, to coordinate interactions between the AI modules. This hybrid approach avoids the current fragmentation between tools designed for healthcare professionals and those designed for patients.
- ***Advanced coordination***: The platform detects delays or gaps in the care pathway, such as unperformed tests or unscheduled appointments, and proposes concrete solutions, such as teleconsultations, secure file transfers or prioritisation at a specialist centre. A shared dashboard provides an overview of the patient journey with progress indicators and contextual alerts, thereby minimising administrative burden while improving continuity of care.

This closed-loop platform is a logical evolution of existing modules and is designed to create synergy between practitioners and patients without ever bypassing clinical judgement.

Detection of atypical data

The detection of atypical data forms the technical basis of the proposals presented in this article. Unlike traditional AI-assisted diagnostic approaches, which aim to identify rare diseases directly, this method focuses on identifying weak signals and unusual patterns within patients' longitudinal data. The objective is to recognise combinations of symptoms, biological changes, or clinical discrepancies that significantly deviate from the trajectories observed in the general population, without ever providing a diagnosis. This approach respects the central role of the clinician, positioning itself as a vigilance-aiding tool rather than a decision-making system. Technically, detection relies on a hybrid architecture that combines several complementary components. An anomaly detection module based on variational autoencoders and temporal graphs analyses structured data (e.g. biological results and ICD codes) and semi-structured data from electronic health records (EHRs). In parallel, a phenotypic matching system uses NLP to standardise clinical notes and align them with the Human Phenotype Ontology (HPO). These two approaches are enriched by patient-reported data and, where available, by measurements from wearable devices. An explanatory large language model (LLM) then generates a comprehensible summary of the detected anomalies, producing traceable chains of reasoning. The entire system operates in near real time while respecting the computational constraints of hospital environments. The system is designed to be highly sensitive to characteristics that are specific to rare diseases, such as the rarity of cases, the temporal variability of symptoms, fragmented data and frequent noise. Rather than classifying patients into specific diseases, it calculates a multidimensional 'atypicality score' that incorporates statistical rarity, temporal consistency, and multimodal consistency of observations. An alert is generated when this score exceeds an adaptive threshold. This design avoids excessive false positives while maintaining high sensitivity to ultra-rare cases where even subtle signals can be decisive.

Optimisation of care pathways

The optimisation of care pathways is the overarching objective of the various proposed AI components. Beyond simply detecting atypical signals or providing informational support, this component aims to streamline the entire patient journey by minimising the administrative, logistical, and time-related obstacles that patients with rare diseases face. This approach does not involve automating medical decisions; rather, it intelligently orchestrates the stages of the pathway to reduce delays, redundancies, and missed opportunities while preserving the central roles of clinicians and patients. This optimisation is based on a longitudinal and systemic view of the healthcare pathway. Technically, optimisation relies on a multi-objective constrained scheduling algorithm that considers

medical constraints (e.g. intervals between tests, biological prerequisites and drug interactions), logistical constraints (e.g. availability of specialist centres, geographical distances and capacity of technical facilities) and patient preferences (e.g. schedules, mobility and cognitive load). An agent-based orchestration engine coordinates various modules in real time, including practitioner alerts, patient assistance, and a knowledge base for rare disease pathways. When delays or inconsistencies are detected (e.g. a test not being carried out within the recommended timeframe), the system proposes realistic alternative solutions (such as a teleconsultation or priority rescheduling) whilst informing the attending doctor, with the patient's consent. This optimisation involves more than just planning. It incorporates the monitoring of treatment adherence and tests, the proactive management of missed appointments and the reduction of redundant tests, all of which are achieved through improved information flow between stakeholders. The patient assistant can provide personalised reminders of upcoming steps in the form of voice or text notifications tailored to the patient's profile, while the practitioner benefits from a comprehensive, up-to-date overview of the care pathway. This creates a virtuous cycle where the early detection of anomalies drives optimisation and where optimisation enhances the relevance of alerts.

Empowering patients

Patient empowerment is an essential and cross-cutting pillar of the proposal. In the context of rare diseases, where patients and their families often have to take on a significant amount of responsibility for coordinating and understanding their care pathway, AI can play a decisive role in strengthening their autonomy and their ability to actively participate in decisions that affect them. The intention is not to replace dialogue with healthcare professionals, but rather to equip patients with the necessary tools to improve their understanding of their situation, enable them to ask pertinent questions, and facilitate their collaboration with practitioners in an informed manner. This approach fully respects the ethical principle that patients are the primary actors in their own care journeys. The multimodal AI assistant directly contributes to this by transforming complex medical information into accessible, personalised content. It explains reports, genetic results and prescriptions in plain language, tailored to each patient's level of health literacy. It also automatically generates diagrams, infographics and short animations to visualise disease mechanisms, treatment considerations and the stages of the care pathway. Thanks to voice and image recognition, patients can photograph a document or ask a question verbally in natural language. The assistant then responds in a contextualised, reassuring, step-by-step manner, regularly checking the patient's understanding. This significantly reduces anxiety caused by lack of understanding and boosts patients' confidence in their ability to manage their condition. In addition to providing explanations, the AI supports task automation by offering practical guidance on daily procedures. The assistant helps to prepare for consultations by suggesting relevant questions, summarising medical histories and anticipating discussion points. They also facilitate coordination by proposing an optimised schedule, reminding patients of tests and gently monitoring adherence. Patients retain full control over the data they share: at each stage, they decide what information to disclose to the doctor. This transparency reinforces the sense of agency and reduces the helplessness often experienced during a long diagnostic journey.

Guidelines for technical solutions

Implementing the presented proposals relies on a combination of mature and emerging technologies that have been selected for their clinical feasibility and compliance with ethical and regulatory standards. At the heart of the system are several complementary building blocks. Natural language processing (NLP) and large language models (LLMs), which have been fine-tuned using medical data, enable us to gain a detailed understanding of unstructured clinical notes and generate explanations that are accessible to patients. These models are complemented by unsupervised machine learning techniques, such as variational autoencoders and temporal graph methods, which are highly effective at identifying anomalies in sparse, longitudinal data. Knowledge graphs incorporating the Orphanet, Human Phenotype Ontology (HPO) and Mondo Disease Ontology ensure the standardisation and interoperability of phenotypic data. Finally, agent-based approaches inspired by systems such as DeepRare orchestrate multi-step reasoning by combining information extraction, phenotypic matching, knowledge base searches, and generation of traceable explanations. This technological stack enables the sensitive detection of atypical signals and the provision of personalised patient support. From an architectural perspective, a hybrid cloud/on-premises approach is favoured. Computationally intensive modules, such as the training and inference of

LLMs (large language models) and anomaly detection models, are hosted on a sovereign cloud infrastructure that is compliant with the Health Data Hub or a certified HDS (Health Data Host) operator. Components that are critical in terms of confidentiality and latency, such as real-time alert generation and the clinician interface, operate on-premises within healthcare facilities. This hybrid approach guarantees performance, scalability, and data sovereignty. The system is designed to meet the stringent requirements of the European AI Regulation (AI Act), which classifies such systems as 'high-risk'. This requires enhanced transparency, detailed technical documentation, mandatory human supervision, and explainability mechanisms such as reasoning chains, attention maps and audit reports. Ultimately, any decision or alert generated by the AI remains under the control of the clinician or patient.

These solutions must be evaluated rigorously and using multiple criteria. A prospective multicentre study conducted across several French reference centres would be a decisive step, although it is clear that implementing such a study would be difficult, if not impossible. Primary criteria would include reducing the average time to diagnosis, decreasing the number of unnecessary consultations and tests, and improving patient and practitioner satisfaction scores (measured via validated questionnaires such as the CSQ-8 or specific scales). Secondary criteria would include the costs saved for the healthcare system. Secondary indicators would focus on acceptability, the cognitive burden on doctors and patients, and patients' health literacy and ability to act independently. An independent assessment of biases (e.g. ethnic, geographical and socio-economic) and robustness in the face of incomplete data could also be carried out.

Finally, the success of the project depends on a strong partnership ecosystem. Rare disease reference centres and the Imagine Institute would provide clinical expertise and access to real-world data. Orphanet will provide up-to-date knowledge bases, while patient organisations such as AFM-Téléthon and EURORDIS will ensure the co-development process is centred on users' actual needs. These collaborations will validate the clinical relevance of the solutions and anchor the project in an ethical and inclusive approach, placing respect for human dignity and reducing inequalities at the heart of technological development.

Discussion

The proposal for an integrated AI system to minimise delays and bottlenecks in care pathways for rare diseases presents several significant challenges, encompassing technical, ethical, and organisational aspects. From a technical perspective, the main challenges lie in managing the rarity of cases, the variability of symptoms over time, and the fragmentation of health data. While hybrid approaches involving anomaly detection, knowledge graphs and AI agents show promising results in controlled environments, their effectiveness in real-world conditions, particularly in the presence of EHR noise and missing data, remains to be demonstrated. Another critical hurdle is the explainability of models: practitioners must be able to understand and validate the alerts generated, which requires the development of advanced traceability mechanisms that comply with the AI Act.

From an ethical and regulatory perspective, the main risk is to the preservation of the patient–doctor relationship. AI must remain a decision support tool, not a substitute. Therefore, granular consent, transparency regarding data use, and the ability for patients to withdraw consent at any time are essential. Potential biases, such as the under-representation of certain ethnic or geographical populations in the training data, could exacerbate inequalities in access to diagnosis. A rigorous assessment of the impact on vulnerable groups will therefore be necessary. From an organisational perspective, acceptance by practitioners is key to success. The system must be non-intrusive and integrated with existing clinical software, and it must not increase cognitive load. Implementation studies must measure clinical effectiveness and the impact on consultation time and practitioner satisfaction. Finally, economic issues such as deployment costs and return on investment, and legal issues such as liability in the event of an error, will need to be clearly addressed.

Despite these challenges, the outlook is encouraging. Well-designed AI has the potential to significantly shorten the diagnostic pathway, improve care coordination and empower patients to take control of their own care journey, all the while respecting the central role of the medical profession.

Conclusion

Rare diseases are still characterised by prolonged diagnostic uncertainty and multiple bottlenecks throughout the care pathway. In

light of this, artificial intelligence is a promising tool for augmenting clinical judgement and actively supporting patients.

This article proposes an integrated framework based on three complementary pillars: an alert system for unusual data for clinicians; a patient-centred, multimodal assistant; and a 'closed-loop' platform linking the two. These solutions are based on realistic technologies, such as natural language processing (NLP), large language models (LLMs), unsupervised machine learning, knowledge graphs and agentic AI, whilst adhering to high ethical and regulatory standards. The proposed approach aims to minimise the main obstacles in the diagnostic journey, such as delayed diagnosis, fragmented information, difficulty understanding medical documents, logistical burdens, and loss of patient autonomy. By placing both practitioners and patients at the centre, it aims to transform the diagnostic odyssey into a smoother, more coordinated, and more humane journey.

Future work should prioritise conducting multicentre, prospective studies to assess the actual impact on time to diagnosis, unnecessary procedures, user satisfaction, and avoided costs. The success of these innovations will depend on close collaboration between reference centres, research institutes, patient organisations and industry, while ensuring transparency, data sovereignty and equitable access. Ultimately, AI should be viewed as an ally of medicine, capable of freeing up precious clinical time and empowering those affected by rare diseases. Responsible implementation could be a significant step towards earlier, more personalised and more inclusive care.

References

1. AlSaad R., et al. "Temporal modelling of irregular clinical trajectories in electronic health records for rare diseases". Artificial Intelligence in Medicine (2024).
2. Boardman-Pretty F., et al. "Evaluating algorithmic approaches to rare disease case-finding: a retrospective validation study using electronic health records". Orphanet Journal of Rare Diseases (2026).
3. Digitale J., et al. "Methods for addressing missingness in electronic health records when building clinical prediction models". JMIR Medical Informatics (2025).
4. Doan LMT., et al. "Bridging modalities with AI: a review of AI advances in multimodal biomedical imaging". Commun Eng 5 (2026): 30.
5. Faviez C., et al. "Enhancing rare disease detection with deep phenotyping from EHR narratives: evaluation on Jeune syndrome". International Journal of Medical Informatics 203 (2025): 106021.
6. Gebeyehu B., et al. "Detection of rare medical events in electronic health records using machine learning: Current practices and suggestions – A scoping review". PLOS ONE (2026).
7. Germain DP. "Applying artificial intelligence to rare diseases: a literature review highlighting lessons from Fabry disease". Orphanet J Rare Dis 20 (2025): 186.
8. Greco KF., et al. "A weakly supervised transformer for rare disease diagnosis and subphenotyping from EHRs with pulmonary case studies". npj Digital Medicine (2026).
9. Groza T, Baynam G and Jamuar SS. "Reimagining care of people living with rare diseases with artificial intelligence". PLOS Medicine 23.2 (2026): e1004966.
10. Kolekar SS., et al. "AI Systems for Diagnosing Rare Diseases: Challenges and Solutions". Conference: ICIMMI 2024: 6th international conference on information management & machine intelligence (2024).
11. Lin S., et al. "An artificial intelligence-based approach for identifying patients with rare diseases from electronic health records: An artificial intelligence-based approach for identifying patients with rare diseases using retrospective electronic health records applied to Pompe disease". Frontiers in Neurology (2023).
12. Mao X., et al. "A phenotype-based AI pipeline outperforms human experts in differentially diagnosing rare diseases using EHRs". npj Digital Medicine (2025).
13. Patharkar A., et al. "Predictive modelling of biomedical temporal data in healthcare applications: review and future directions". Front. Physiol 15 (2024).
14. Salmi M., et al. "Handling imbalanced medical datasets: review of a decade of research". Artificial Intelligence Review 57 (2024):

273.
15. Seifi S., et al. “Complying with the EU AI Act: Innovations in explainable and user-centric hand gesture recognition”. Machine Learning with Applications 20 (2025): 100655.
16. Shyr C., et al. “Large Language Models for Rare Disease Diagnosis at the Undiagnosed Diseases Network”. JAMA Network Open 8.8 (2025): e2528538.
17. Sreenivasan M., et al. “Temporal and comorbidity-aware representation of longitudinal patient trajectories from electronic health records”. Biomedical Physics & Engineering Express 12.1 (2026).
18. Visibelli A., et al. “The Impact of Artificial Intelligence in the Odyssey of Rare Diseases”. Journal of Personalized Medicine (2023).
19. Viswan V., et al. “Multimodal fusion and explainability of artificial intelligence models in Alzheimer’s Disease detection”. Brain Informatics 13.1 (2026): 5.
20. Xin Wang, Da He and Chunlin Jin. “AI-driven enhancements in rare disease diagnosis and support system optimization”. Intractable Rare Dis Res 14.4 (2025): 306-308.
21. Magda Wojtara., et al. “Artificial intelligence in rare disease diagnosis and treatment”. Clin Transl Sci 16.11 (2023): 2106-2111.
22. Zhao W., et al. “An agentic system for rare disease diagnosis with traceable reasoning”. Nature (2026).
23. Zhu M., et al. “Bridging data gaps of rare conditions in ICU: a multi-task domain adaptation approach”. npj Digital Medicine (2026).

Foot Notes

1. The General Data Protection Regulation.
2. Health Insurance Portability and Accountability 1996 Act.
3. Electronic Health Record.
4. International Classification of Diseases, 11th Revision (ICD-11) – the WHO’s standard for describing and comparing rare diseases.
5. Digital Imaging and Communications in Medicine.
6. Fast Healthcare Interoperability Resources: a data transfer format designed to improve the interoperability of healthcare systems by facilitating the exchange of medical and administrative data between stakeholders in the healthcare ecosystem.