prompt opinion™

WHITE PAPER

The Enrollment Gap

Why eligible clinical trial participants are already in your EHR

How academic medical centers can unlock research enrollment from the patients they already serve, using AI-powered longitudinal patient intelligence.

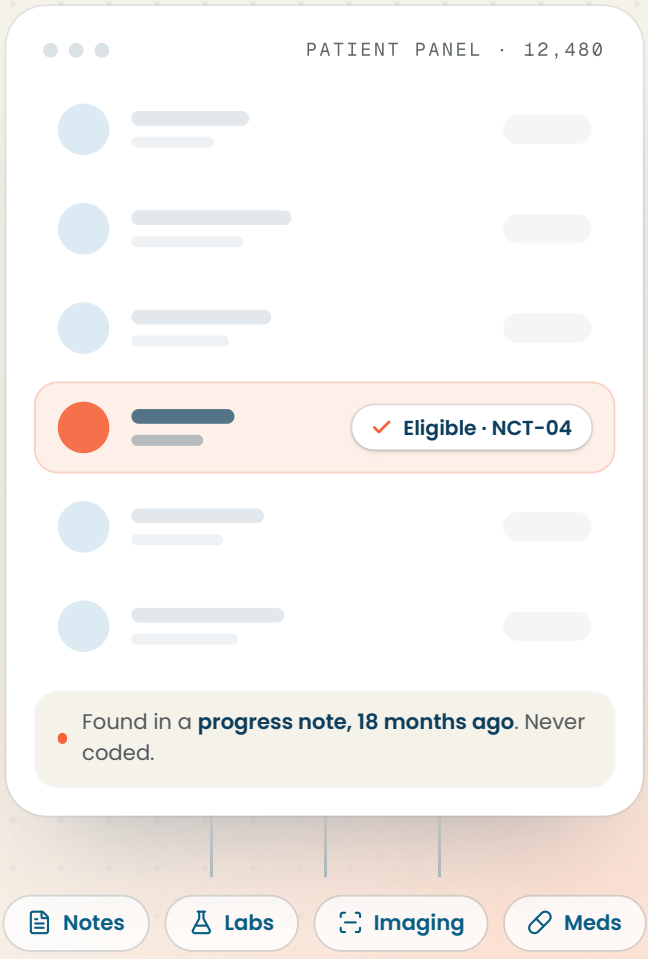

AI-powered longitudinal patient intelligence

September 2026 · 18 min read

promptopinion.ai

EXECUTIVE SUMMARY

UP TO $8M/day

lost for every day a late-stage trial runs behind

~80%

of clinically relevant patient data sits in unstructured records

Academic medical centers run some of the most complex and consequential clinical research in the world. Yet a persistent and largely invisible problem undermines their ability to enroll patients into sponsored and investigator-initiated trials. The patients who qualify are already inside their electronic health records, and they are not being found.

Approximately 80% of clinical trials fail to meet their original enrollment timelines, and patient recruitment is the single most cited reason. The consequences extend beyond the sponsoring pharmaceutical company. For academic medical centers, missed enrollment means missed research funding, reduced investigator productivity, diminished institutional reputation in competitive grant environments and, critically, patients who never received access to potentially life-changing therapies.

The root cause is structural. Up to 80% of clinically relevant patient information resides in unstructured data: physician notes, discharge summaries, imaging reports and free-text narratives that traditional EHR search tools cannot systematically query. Research coordinators, already stretched across multiple active studies, spend a disproportionate share of their time on manual chart review, often reviewing dozens of records to identify a single eligible participant.

> **The enrollment gap is not primarily a population problem. It is a data access problem. The patients exist. The data exists. What is missing is the intelligence layer that connects them to active research protocols continuously, accurately and without adding burden to investigators or research operations teams.**

Prompt Opinion addresses this gap through AI-powered longitudinal patient intelligence: a FHIR-native platform that ingests, structures and continuously queries both structured and unstructured EHR data against active trial protocols. It delivers ranked, pre-screened patient lists to research coordinators rather than requiring them to generate those lists manually. The result is more studies enrolling faster, fewer coordinators stretched across an unsustainable pre-screening burden, and a measurable improvement in sponsored research revenue from the patient population the institution already serves, without proportional increases in research operations cost.

01 The scale of the enrollment problem

Clinical trial enrollment is in a state of systemic underperformance. According to data from the Tufts Center for the Study of Drug Development, approximately 80% of clinical trials fail to meet their original enrollment timelines.[1] A CenterWatch study of several hundred investigative sites found that upwards of 90% of clinical trials fail to meet timelines as indicated in the study contract,[2] and that 11% of research sites fail to enroll a single patient in a given trial.[3]

The consequences fall on multiple stakeholders. For pharmaceutical sponsors, each day of delay in a late-stage oncology trial can cost between $600,000 and $8 million in foregone revenue.[4] Across the industry, slow enrollment is estimated to consume $50 billion annually in direct costs, extended CRO contracts and the downstream value destruction of delayed launches.[1]

90%[2]	**11%**[3]	**$50B**[1]
of trials miss contracted timelines	of sites enroll zero patients	lost to slow enrollment yearly
1 in 5[5]	**12.2 mo**[5]	**31%**[5]
trials finish on planned time	median completion delay	median enrollment shortfall

Enrollment performance, by the numbers.

A cross-sectional study published in JAMA Surgery, examining 2,542 randomized clinical trials registered on ClinicalTrials.gov between 2010 and 2014, found that approximately 1 in 5 trials were completed within their planned timeframe, with a median delay of 12.2 months. Less than half of trials met their prespecified enrollment target upon completion, with the median shortfall accounting for 31% of the planned sample size.[5]

For academic medical centers, the impact is less visible but equally consequential

- Sponsored research funding that was budgeted for a trial does not materialize if enrollment targets are not met, directly affecting research revenue.
- Investigators whose studies fail to enroll become less competitive for future grant funding and industry partnerships.
- Patients who could have accessed novel therapies through a clinical trial are never identified or approached.

- Institutional reputation in competitive research networks is shaped, in part, by a site's track record of delivering enrolled patients on schedule.

> **These are not the failures of individual investigators or research teams. They reflect a structural problem in how eligible patients are identified, one that technology has not yet adequately solved.**

02 Why the problem persists: the data accessibility gap

The conventional explanation for poor enrollment is that there are not enough eligible patients. In most cases, this is inaccurate. The more precise explanation is that eligible patients exist within health system EHRs but cannot be systematically identified using current tools.

2.1 The unstructured data problem

Up to 80% of clinically relevant information in electronic health records resides in unstructured data: physician notes, discharge summaries, radiology reports, operative notes and free-text clinical narratives.[6] This is not a marginal share of the clinical record. It is the majority of it. And critically, it is the portion most likely to contain the nuanced clinical detail, such as disease severity, prior treatment response and comorbidity specifics, that determines trial eligibility.

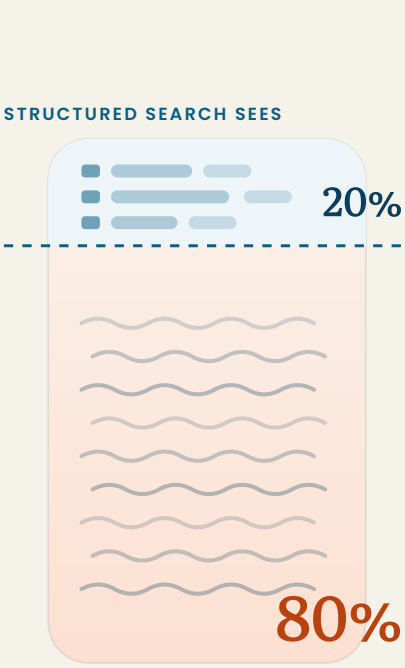

STRUCTURED FIELDS 20%

Diagnosis codes · Medications · Lab values · Demographics

UNSTRUCTURED RECORD 80%[6]

Physician notes · Discharge summaries · Radiology reports · Operative notes · Free-text narratives

Where disease severity, prior treatment response and comorbidity detail actually live.

Most of the record sits below the waterline of structured search.

Research published in the Journal of the American Medical Informatics Association demonstrated that structured EHR data achieved an average recall of 51.7% in identifying clinical concepts relevant to patient eligibility, while analysis of unstructured notes achieved 95.5% recall for the

same task.[7] In practical terms, a recruitment process that relies only on structured EHR fields will miss approximately half of the patients who actually qualify.

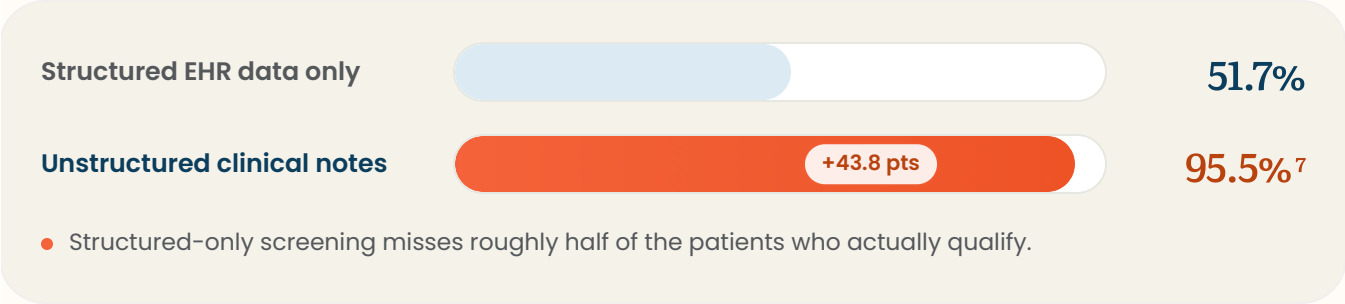

Recall in identifying eligibility concepts. Structured-only screening misses roughly half of the patients who qualify.

A study published in medRxiv analyzing over 19 million EHR notes from more than 500,000 patients found that NLP analysis of unstructured notes contributed a 29.5% relative increase in patients identified as smokers beyond what structured data alone captured, and a 19.3% relative increase for obesity.[8] For eligibility criteria that depend on detailed clinical history, as most trial protocols do, the gap between structured-only and full-record analysis is material.

2.2 The protocol complexity problem

Clinical trial inclusion and exclusion criteria are written in clinical language: free text, medical terminology and context-dependent conditions that cannot be directly translated into structured database queries. A coordinator screening patients for a complex oncology or rare disease trial may be working against 30 to 50 individual criteria, each requiring a different type of data from a different part of the patient record.

Tufts CSDD studies have consistently shown a positive association between protocol complexity, longer timelines, poorer patient recruitment and retention rates, and a higher incidence of protocol amendments.[2] More than 40% of trials amend their protocol before the first subject visit, delaying trials by an average of four months.[9]

2.3 The coordinator burden problem

The current standard for patient identification in most research operations is manual chart review by clinical research coordinators. A time-and-motion study of clinical research coordinators in a pediatric emergency department found that patient screening consumed 31.6% of coordinator time, the single largest category of their work.[10]

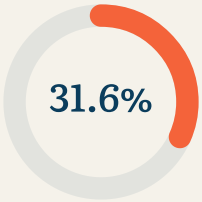

of coordinator time is spent screening patients[10]

The single largest category of coordinator work, measured in a pediatric emergency department time-and-motion study. At an institution running many concurrent studies, it multiplies.

- Screening is the single largest category of coordinator work.

An information extraction study examining heart failure patient screening found that an experienced research coordinator required two weeks, 80 hours, to manually pre-screen 198 patient records against a single trial's eligibility criteria, ultimately identifying 40 potentially eligible patients for further review.[11] An automated AI system completed the same pre-screening task in approximately 20 minutes.[11]

Manual review — **80 hours**

two weeks, one research coordinator · each block is one working day

Automated — **~20 minutes**[11]

same task, same 198 records

SAME RESULT

40

eligible patients surfaced

- Pre-screening 198 records against one protocol: two weeks by hand, twenty minutes automated.

A commentary published by the Association of Clinical Research Professionals (ACRP) in 2026 noted that research staff currently spend hours on pre-screening and must often manually evaluate large pools of potential participants for trial eligibility, diverting time and energy away from meaningful patient interaction.[12] At an institution running multiple concurrent studies, this burden is multiplicative, and it compounds directly into enrollment underperformance.

03 What longitudinal patient intelligence changes

The enrollment gap is not closed by recruiting harder. It is closed by finding patients more accurately and earlier, using the data that already exists within the institution's EHR. This requires a fundamentally different approach to how patient records are structured, queried and connected to active research protocols.

3.1 From episodic to longitudinal: the core distinction

Most EHR search tools operate on episodic data: what is recorded in structured fields during a specific encounter. A longitudinal patient record, by contrast, synthesizes the full clinical history of an individual patient across all encounters, data types and time periods. Structured fields, unstructured notes, lab values, imaging reports, medication history and historical visits, in one place.

This distinction is material for clinical trial matching. Many eligibility criteria require evidence of a condition or prior treatment over a defined historical period, not just what was coded in the last encounter. A patient who received a specific therapy 18 months ago, documented only in a clinical note, is invisible to structured-field search. They are not invisible to longitudinal patient intelligence.

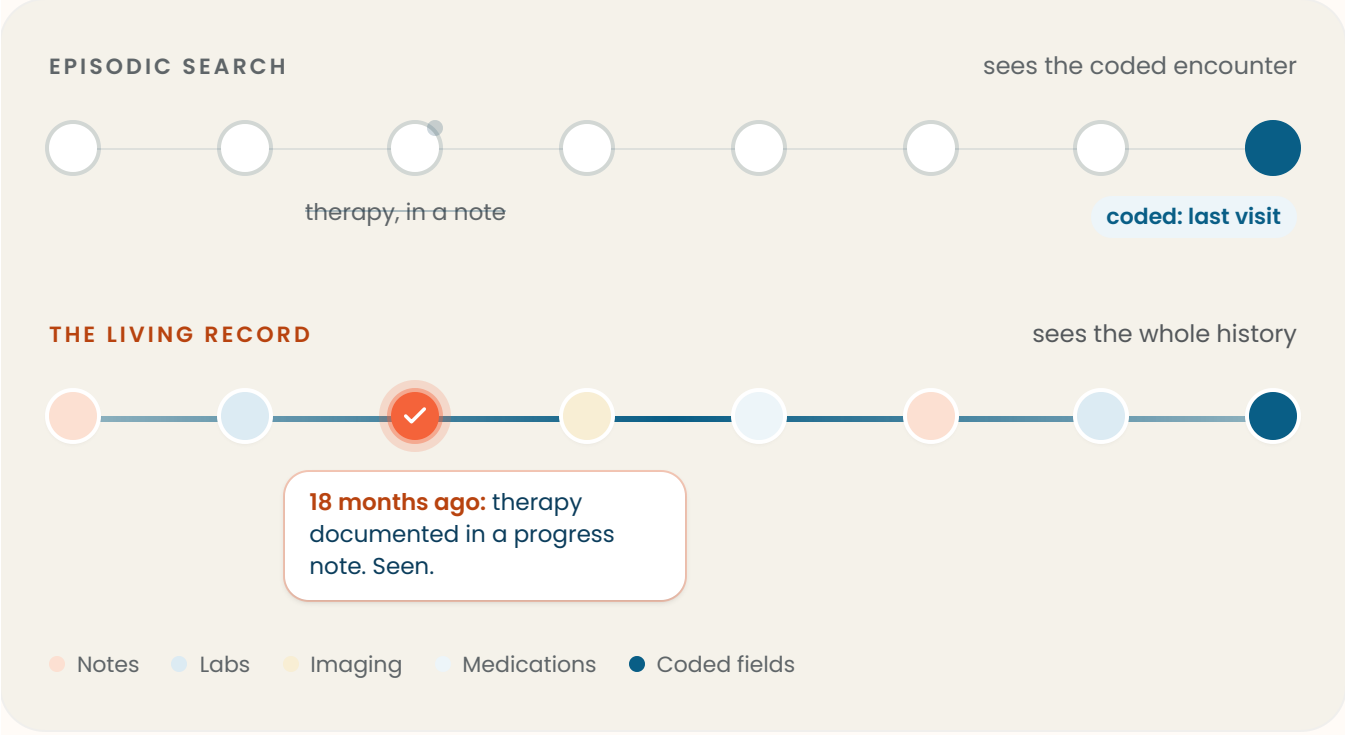

Episodic search sees the coded encounter. The living record sees the whole history, notes included.

3.2 How Prompt Opinion works

Prompt Opinion is an AI-powered clinical intelligence platform built natively on FHIR (Fast Healthcare Interoperability Resources) standards. Developed in close collaboration with health system partners, the platform ingests both structured and unstructured EHR data across a broad range of clinical data types, including clinical notes, laboratory results, radiology reports, pathology, medications, procedures, encounters and genomics, and makes this data continuously queryable against active research protocols.

The platform connects to health system EHR environments, including Epic, through SMART on FHIR APIs: the same standardized, vendor-approved integration framework that major EHR systems already expose for third-party applications. Implementation works within existing data infrastructure rather than requiring custom integrations, parallel data pipelines or direct database access.

For health IT and security teams evaluating the platform, Prompt Opinion operates inside a certified environment: HITRUST e1, SOC 2, HIPAA and ONC Certified Health IT. These are the core security and compliance frameworks that health system security and vendor assessment teams require, and they are available for review as part of any institutional security assessment.

The clinical trial matching workflow

1. **Protocol ingestion**
   Inclusion and exclusion criteria are extracted from trial protocols and translated from clinical language into structured, queryable logic.

2. **Longitudinal record construction**
   The platform builds a continuously updated, FHIR-native patient record by ingesting structured EHR fields, clinical notes, lab reports, imaging and historical encounter data, normalizing terminology across ICD, SNOMED, LOINC and RxNorm and reconciling events across encounters and time.

3. **AI eligibility matching**
   Each patient's longitudinal record is matched against the extracted protocol criteria, with the AI scoring and ranking patients by likelihood of eligibility.

4. **Missing data identification**
   For patients who are close to eligible but have a data gap, for example a required lab value not yet on record, the system flags them separately so coordinators can decide whether the gap can be resolved, rather than dismissing a potentially eligible patient.

5 Coordinator dashboard

Coordinators receive a prioritized, actionable worklist rather than a chart review queue: new matches, candidates needing a specific lab review, patients ready for outreach and cases requiring clinical review. The coordinator's role shifts from searching to acting. Each candidate arrives with the matching rationale surfaced and the next step pre-determined.

This process runs continuously, not as a one-time chart pull. As new patients enter the system and existing records are updated, eligibility is re-evaluated automatically across every active protocol at once. An institution running multiple concurrent studies does not need parallel chart review processes for each.

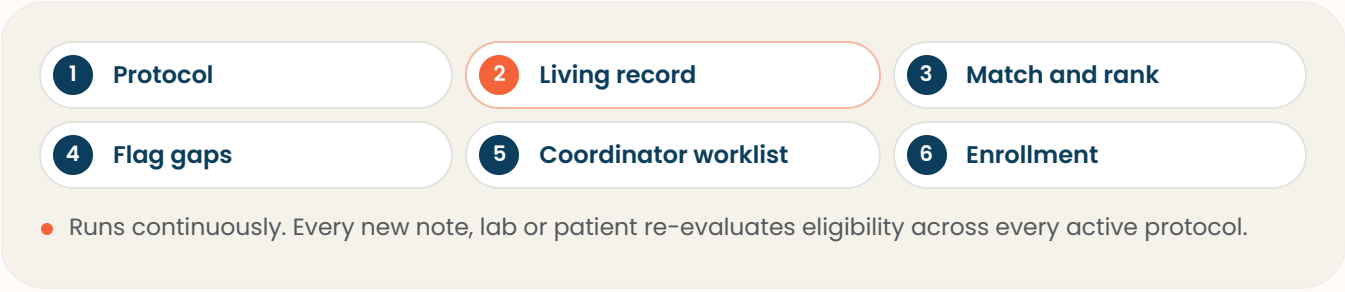

From protocol to enrolled patient, one continuous loop.

3.3 A note on implementation and burden

Any new technology introduced into a clinical research environment must be evaluated not only on its potential benefit but on the burden it places on already-stretched teams. Prompt Opinion is designed with this constraint as a first principle, not an afterthought.

For research coordinators

The platform reduces workload rather than adding to it. Coordinators no longer generate a list of potentially eligible patients. They receive one, pre-screened, ranked and with next actions identified. Their time moves from low-value pre-screening to patient interaction, consent conversations and study management.

For investigators and physicians

Investigators are not asked to change their clinical workflows. Patient identification happens outside the clinical encounter, using existing EHR data. The first point of investigator contact is a pre-screened, ranked candidate list, not a request to conduct a chart review. Cases needing clinical review surface only when human judgment is required.

For health IT and implementation teams

The FHIR-native, SMART on FHIR integration approach means the platform operates within the data access frameworks that health systems and EHR vendors have already established and approved. HITRUST e1, SOC 2, HIPAA and ONC certification mean the security review starts from a strong baseline, reducing the assessment burden on internal teams.

04 Institutional impact: research operations and research revenue

For academic medical centers, the consequences of systematic enrollment underperformance fall across two dimensions that are deeply connected but rarely addressed together: the operational burden on research teams, and the revenue yield from the institution's sponsored research portfolio. Improving enrollment performance addresses both simultaneously.

4.1 Research operations

Systematic enrollment underperformance creates a specific set of operational and strategic pressures for research leadership:

- Studies that miss enrollment targets require timeline extensions, additional site activations or protocol amendments, each of which consumes research operations resources without proportional benefit.
- Coordinators allocated to studies that enroll poorly are effectively under-utilized relative to their cost, while coordinators on high-demand studies are over-stretched.
- Investigator confidence in the institution's ability to deliver enrollment affects which studies investigators choose to pursue and whether external sponsors select the institution as a site.
- Research networks and multi-site collaborations increasingly track per-site enrollment performance. Poor performance affects future inclusion.

AI-assisted patient matching addresses each of these pressures by improving the yield of the existing research operations infrastructure: more enrolled patients per coordinator hour, more studies meeting their enrollment targets, and a stronger institutional track record for future sponsor engagement.

> **The system does not selectively improve enrollment for a single study in isolation. Because it runs continuously across all active protocols simultaneously, the benefit compounds across the institution's full research portfolio.**

4.2 Research revenue

The financial case for improving clinical trial enrollment at an academic medical center operates on two dimensions: revenue generation and cost efficiency.

On the revenue side

Sponsored clinical trials generate direct revenue through per-patient reimbursement from sponsors, overhead recovery on research grants and indirect cost recovery on federally funded studies. When enrollment targets are not met, a portion of this budgeted revenue does not materialize. The magnitude depends on portfolio size, trial complexity and average per-patient sponsor budgets, but for a major academic medical center running a substantial number of sponsored trials annually, the cumulative impact of systematic under-enrollment is meaningful.

On the cost efficiency side

Research coordinator time is the primary variable cost in clinical trial operations. When coordinators spend the majority of their time on manual pre-screening, cost per enrolled patient is high and throughput per coordinator is low. Reducing the pre-screening burden through automation does not require adding coordinator headcount to enroll more patients. It improves the yield of existing headcount. Revenue grows with enrollment, while the marginal cost of each additional enrolled patient falls as automation scales.

A validation conducted with a major academic medical center demonstrated strong clinical interest and feasibility, with institutional security review processes, standard for health systems of this scale, identified as the primary pathway to full deployment. This validation confirms both the technical capability of the platform and the appetite among research leadership for this category of solution.

05 Looking forward: from matching to research intelligence

The immediate application of AI-powered longitudinal patient intelligence is clinical trial matching: identifying eligible participants faster, with less coordinator burden, across a broader share of the EHR population. But the infrastructure required to do this well, a continuously updated, FHIR-native, AI-queryable patient record, creates the foundation for a broader set of capabilities that academic medical centers will increasingly need.

Site feasibility modeling

Real patient population data, rather than coordinator estimates, to evaluate whether a proposed trial is feasible at a site before it is activated.

Protocol feasibility analysis

Draft inclusion and exclusion criteria queried against the population before a protocol is finalized, surfacing criteria that will constrain recruitment and reducing costly mid-trial amendments.

Research dashboards

Real-time visibility for research operations leadership into enrollment status, screening yield and coordinator activity across all active studies.

Registry creation

Continuously updated patient registries for specific conditions or biomarker profiles, enabling faster activation of future studies.

Investigator copilots

Relevant patient cases and research opportunities surfaced to investigators inside their existing clinical workflows, without a separate research system to query.

- Once the living record is in place, the same infrastructure supports the next five.

The path from today's enrollment gap to tomorrow's research intelligence infrastructure runs through a single decision: making the patient data that already exists inside the EHR continuously accessible, structured and queryable for research purposes. Prompt Opinion is built to make that transition as low-friction as possible for institutions ready to take it.

CONCLUSION

The clinical trial enrollment problem is well understood. Its persistence is not a reflection of insufficient effort. It is a reflection of a structural gap between the data that exists in health system EHRs and the tools available to translate that data into research participation.

For academic medical centers, closing this gap is not merely an operational improvement. It is a strategic imperative: for research revenue, for investigator retention, for institutional competitiveness in sponsored research, and for fulfilling the fundamental mission of translating clinical knowledge into patient benefit.

> **The patients who qualify for your active studies are already in your EHR. The question is whether your current infrastructure can find them.**

NEXT STEP

The patients are already in your EHR.

See how the living record finds them, ranked and cited, for the studies you are running today.

promptopinion.ai/solutions/clinical-trials · promptopinion.ai/contact

APPENDIX A

Platform architecture

The four-layer architecture of the Prompt Opinion platform, for review by health IT, informatics and security teams evaluating technical requirements.

LAYER 1

Health system data

Everything the institution already holds, connected via SMART on FHIR.

EHR · Notes · Labs · Radiology · Pathology · Medications · Procedures · Encounters · Genomics · Registries

LAYER 2

Longitudinal patient intelligence

The living record: one continuously updated, cited timeline per patient.

- Extracts structured and unstructured clinical concepts
- Maps patient information into a FHIR-aligned model
- Builds a continuously updated patient timeline
- Identifies missing and conflicting clinical information
- Normalizes terminology across ICD, SNOMED, LOINC, RxNorm
- Reconciles events across encounters and time
- Preserves source evidence and provenance
- Connects via SMART on FHIR APIs, compatible with Epic and major EHR platforms

LAYER 3

Clinical research intelligence

Protocols become queries. Patients become ranked, explained candidates.

Protocol intelligence · Patient discovery · Eligibility reasoning · Missing evidence detection · Coordinator dashboard · Recruitment pipeline · Research analytics

LAYER 4

Research operations

Where the coordinator's day now starts.

Ranked candidates · Coordinator worklists · Clinical review · Patient outreach · Consent · Enrollment · Portfolio metrics

CERTIFIED ENVIRONMENT — HITRUST e1 · SOC 2 · HIPAA · ONC Certified Health IT

REFERENCES

1 **Vision Life Sciences Clinical Team.** "Clinical Trial Patient Enrollment: Strategies 2026." Vision Life Sciences, February 28, 2026. visionlifesciences.com/insights/clinical-trial-patient-enrollment-strategies

2 **CenterWatch.** "Enrollment Performance: Weighing the 'Facts'." Applied Clinical Trials Online. appliedclinicaltrialsonline.com/view/enrollment-performance-weighing-facts

3 **Antidote.** "25+ Useful Clinical Trial Recruitment Statistics for Better Results." Antidote, January 19, 2026. antidote.me/blog/25-useful-clinical-trial-recruitment-statistics-for-better-results

4 **Getz KA.** "Evaluating the Impact of Patient Recruitment on Clinical Trial Delays." Tufts Center for the Study of Drug Development, cited in Online Patient Recruitment in Clinical Trials, NCBI PMC7673977. ncbi.nlm.nih.gov/pmc/articles/PMC7673977

5 **McCulloch P, et al..** "Analysis of Rates of Completion, Delays, and Participant Recruitment in Randomized Clinical Trials in Surgery." JAMA Surgery, PMC9857498. ncbi.nlm.nih.gov/pmc/articles/PMC9857498

6 **Verana Health.** "Extracting Value from Electronic Health Records by Processing Unstructured Data." February 13, 2024. veranahealth.com/extracting-value-from-electronic-health-records-by-processing-unstructured-data

7 **Hernandez-Boussard T, et al..** "Unstructured EHR Data More Useful for Predictive Analytics, Study Shows." Healthcare IT News, October 31, 2019. healthcareitnews.com/news/unstructured-ehr-data-more-useful-predictive-analytics-study-shows

8 **medRxiv.** "Using Natural Language Processing of Clinical Notes to Supplement Structured Electronic Health Record Data for Phenotyping Smoking and Obesity in a Healthcare System." January 2026. medrxiv.org/content/10.64898/2026.01.18.26344356.full.pdf

9 **Array Live.** "Mitigating Clinical Trial Enrollment Challenges with Information and Communication." February 8, 2023. arraylive.com/blog/how-to-increase-clinical-trial-enrollment-with-information-and-communication

10 **Tsang R, et al..** "A Time-and-Motion Study of Clinical Trial Eligibility Screening in a Pediatric Emergency Department." PMC6445787. pmc.ncbi.nlm.nih.gov/articles/PMC6445787

11 **arXiv.** "An Information Extraction Approach to Prescreen Heart Failure Patients for Clinical Trials." 2016. arxiv.org/pdf/1609.01594

12 **Association of Clinical Research Professionals (ACRP).** "Reimagining the Clinical Research Coordinator in the Age of AI: A Commentary." June 16, 2026. acrpnet.org/2026/06/16/reimagining-the-clinical-research-coordinator-in-the-age-of-ai-a-commentary