



AI-Powered Patient Matching for Clinical Trials Enrollment

How Intuceo engineered a LangGraph-orchestrated agentic AI system that navigates complex eligibility criteria and fragmented EHR data to accelerate clinical trial recruitment.

The Challenge

Clinical trials are the backbone of drug development. They are also one of the most expensive and frequently delayed phases in the pharmaceutical pipeline. Patient recruitment sits at the center of this problem. Sponsors spend substantial time and money trying to identify eligible subjects, and yet enrollment shortfalls remain the leading cause of trial delays and failures.

For healthcare providers managing millions of cancer patients across distributed systems, the scale of the problem compounds quickly. Matching a patient to a specific clinical trial requires interpreting highly specific eligibility criteria, integrating data across multiple formats and source systems, and doing so at speed and at scale.

\$46B+

Annual clinical trials market

50%

Of trials delayed by enrollment gaps

37%

Of sites miss recruitment target

50%

Dev time reduction with agentic AI

Contact Us



+1 904 404 7520



info@intuceo.com



Florida, U.S

4110 Southpoint Blvd,
Suite 124 Jacksonville, Florida – 32216, USA.

Why Traditional Recruitment Methods Fall Short

Clinical trials are the backbone of drug development. They are also one of the most expensive and frequently delayed phases in the pharmaceutical pipeline. Patient recruitment sits at the center of this problem. Sponsors spend substantial time and money trying to identify eligible subjects, and yet enrollment shortfalls remain the leading cause of trial delays and failures.

— Strict eligibility criteria:

Inclusion and exclusion criteria are highly specific and written in complex clinical language. A single trial protocol may contain dozens of interlocking conditions joined by logical connectors, requiring precise interpretation.

— Fragmented patient data:

Patient information is dispersed across EHRs, lab systems, imaging platforms, clinical notes, and care plan databases. No single system holds a complete picture of a patient's eligibility.

— Recruitment process inefficiencies:

Manual review processes are time-intensive, inconsistent across sites, and difficult to scale. Sites frequently operate without a structured workflow for eligibility screening.

— Awareness gaps:

Both patients and providers frequently lack visibility into active clinical trials relevant to a patient's condition. Even when trials exist, the path from awareness to enrollment is rarely straightforward.

— Logistical barriers:

Trial location requirements, scheduling constraints, and patient availability create friction that further depresses enrollment rates.

— Trust and compliance:

Recruitment efforts must navigate patient trust concerns and regulatory compliance requirements, adding another layer of complexity to an already fragmented process.

A Representative Example from Practice

Study: Phase II Advanced Non-Small Cell Lung Cancer

Eligibility criterion excerpt "Participants with non-squamous or squamous histology NSCLC with stage IIIB or stage IIIC disease who are not candidates for surgical resection or definitive chemoradiation per investigator assessment, or stage IV (metastatic) disease who received no prior systemic treatment for recurrent or metastatic NSCLC."

Matching a real-world patient population to criteria of this complexity requires structured parsing, entity resolution, and clinical reasoning. Manual methods produce inconsistent results. Rule-based systems break down when criteria involve conditional logic and multi-concept entities.

The Solution

Intuceo designed and built an agentic AI system specifically for clinical trial patient matching. The system uses a LangGraph-orchestrated multi-agent workflow to automate the full pipeline: from parsing unstructured study protocols, to resolving clinical entities against medical terminologies, to executing a hybrid matching engine that combines structured query logic with large language model (LLM) reasoning.

The architecture is built to handle the real-world messiness of clinical data. It does not require clean, pre-formatted inputs. Instead, it processes raw ClinicalTrials.gov data and unstructured EHR records to surface eligible patients with precision and auditability.

System Architecture Overview

Query Parser Agent

- Fetches unstructured study data from ClinicalTrials.gov API
- Applies Natural Language Processing for intent recognition
- Converts raw trial text into structured JSON with attributes, entities, and values
- Validates and structures queries before downstream processing

Data Retrieval Agent

- Interfaces with Azure MCP for secure patient data access
- Implements Secure Data Access Protocol (SDAP)
- Aggregates patient records from multiple source systems
- Returns normalized patient data for matching

Matching Agent

- Executes hybrid SQL + LLM matching logic
- Applies AI-powered eligibility assessment against parsed criteria
- Generates confidence scores for each patient-trial pairing
- Surfaces matched patients with supporting reasoning

Result Synthesizer

- Formats output for clinical researcher consumption
- Generates longitudinal patient summaries
- Produces ranked, prioritized match results
- Delivers structured reports to the Matching Results Dashboard

Entity Resolution: Bridging Clinical Language and Structured Data

One of the most technically demanding components of the solution is entity resolution. Clinical trial protocols use natural language to describe conditions, procedures, and patient characteristics. EHR systems store equivalent information as coded data using vocabularies like SNOMED CT. Aligning these two representations is essential for accurate matching and historically requires extensive manual curation.

Intuceo solved this by deploying a GPT-based Agent that maps extracted entities to SNOMED CT codes via direct API interaction. This eliminates the brittleness of rule-based terminology mapping and enables the system to handle novel entity combinations without requiring manual coding updates. The result is a dynamic, self-adapting entity resolution layer that bridges trial criteria and EHR data at scale.

AI Matching Engine: Structured and Semantic Matching Combined

The core matching engine runs two parallel tracks. Structured matching executes rule-based SQL query templates against demographic data and quantifiable lab values, where precision and determinism are essential. Semantic matching applies vector similarity search and LLM-driven analysis across clinical notes and unstructured patient records, where context and nuance matter.

The engine does not simply return a yes or no answer. For each candidate patient, it generates a longitudinal summary, explains the basis for the match, and provides a confidence score. Clinical researchers reviewing results can understand exactly why a patient was surfaced, which supports both operational trust and regulatory auditability.

Technical Stack

The solution is built on a modern, cloud-native architecture designed for regulated healthcare environments. Each component was selected to support compliance requirements, data security, and the orchestration demands of a multi-agent agentic workflow.

Component	Technology / Description
Orchestration Framework	LangGraph - manages agent graph execution, state transitions, and inter-agent communication across the multi-agent workflow
Programming Language	Python - primary language for agent logic, data processing pipelines, and API integrations
Language Model Layer	LLM (Commercial / Local) - supports both cloud-hosted commercial models and locally deployed models for air-gapped or compliance-restricted environments
Database	Microsoft SQL Server (MSSQL) - structured patient data storage, SQL query template execution for deterministic matching
AI & Cognitive Services	Azure OpenAI - GPT models for entity resolution, semantic matching, and longitudinal patient summary generation
Cloud Data Layer	Azure MCP (Model Context Protocol) - secure patient record access, multi-source data aggregation, HIPAA-compliant data handling
Terminology & Coding	SNOMED CT API - real-time entity-to-code resolution for clinical terminology standardization
Trial Data Source	ClinicalTrials.gov API - structured and unstructured trial protocol ingestion for automatic criteria parsing

Security and Compliance

The Azure MCP Data Layer is architected to meet the requirements of healthcare data environments. Patient record access enforces role-based access controls, all data in transit and at rest is encrypted, and audit logging captures every data retrieval and processing event for compliance review. The system is designed to operate within HIPAA-governed data environments without requiring changes to existing security infrastructure.

Benefits and Outcomes

- 50% reduction in application development time. Reusable agentic components and LangGraph's structured orchestration layer significantly accelerate the build cycle compared to custom, point-solution development approaches.
- Higher patient matching accuracy against inclusion and exclusion criteria. The hybrid engine resolves entity terminology gaps and handles multi-concept eligibility conditions that rule-based systems routinely fail to process correctly.
- Matching rationale generated for every result. Clinical researchers receive a confidence score and a plain-language explanation for each candidate, supporting faster review decisions and auditability in regulated environments.

Ready to accelerate your clinical trial enrollment?

Intuceo's PhD-led engineering team applies the same rigorous, data-architecture-first approach to your clinical operations challenges.

[Request a Strategic Consultation](#)

