

UNLOCKING CLINICAL TRIAL ENROLLMENT WITH AI SCREENING

It's not a patient population problem.
It's a signal-extraction problem.

By Sunil Gupta, MD, and Joshua Koo, MS, BS



Every practice that runs clinical trials knows a particular kind of frustration. You open a new protocol for geographic atrophy, wet AMD, or diabetic macular edema, confident that patient volume will fill it. And then you watch enrollment crawl. The reflex is to assume the shortfall is a supply problem: not enough patients. However, in a busy retina clinic, the eligible patients are there—the problem isn't patient scarcity,

it's patient identification.

On any given clinic day, dozens of patients have the conditions our open protocols were written to study, and the information needed to identify them as trial candidates is already in the record: acuity measurements, medical and treatment histories, laterality notes, and OCT scans. However, we still need a dependable way to identify that information quickly for each patient.

HOW WE SCREEN TODAY

For most sites, pre-screening follows one of two methods, neither of which scales. The first is physician memory or a cheat sheet. The treating specialist must keep the inclusion and exclusion criteria for every active protocol in mind and match against them in real time.

The second method is a manual pre-screener, usually a clinical trial coordinator who reviews charts ahead of


clinic. This is more systematic but labor intensive, limited in reach, and vulnerable to fatigue. A coordinator can review part of the schedule carefully, but patients may still go unreviewed. Both approaches, even when staffed by capable people, have the same result: Eligible patients are missed, screen failure rates stay high, and enrollment lags.

AN AUTOMATED APPROACH

Now, an agentic AI can create a structured screening process that runs in the background for every scheduled

KEY TAKEAWAYS

- ▶ Current clinical trial screening methods rely on physician memory or a manual pre-screener, often leading to missed eligible patients, high screen failure rates, and lagging trial enrollment.
- ▶ An investigational AI model can automatically flag eligible candidates, allowing the physician to initiate the trial conversation and potential formal screening on the same day.
- ▶ Using the AI system allowed the Retina Specialty Institute to bring a protocol live in about 4 weeks.



Counts shown are illustrative and for demonstration only.

Figure 1. This diagram illustrates the five-stage AI screening funnel that runs automatically in the background for every scheduled patient. Counts shown are illustrative. Abbreviation: LLM, large-language model.

patient. The framework we have been working with organizes screening into five stages (Figure 1).

Stage 1: Intake

Every patient on the day's schedule enters the screening process automatically, and the work begins before the clinic doors open.

Stage 2: Extraction Filter

A broad extraction filter casts a deliberately wide net, comparing structured data such as acuity ranges and prior treatment history against protocol criteria. This stage is fast, rule based, and consistent, built to keep candidates in rather than screen them out prematurely.

Stage 3: AI Deep Filter

The third stage is where the real value lies. A large-language model reads the free text of the clinical note and pulls out diagnosis history, prior treatments, laterality, and disease duration, then maps that language against each protocol's inclusion, exclusion, and study eye requirements. Anything ambiguous is flagged for physician review. In addition, image-analysis models quantify pathology from OCT and retinal images, measuring geographic atrophy lesion size, fluid, and layer integrity, and returning objective values to check against protocol thresholds.

Many such models now exist, and they take much of the subjectivity out of the eligibility question. The system reads every chart and measures every image, for every patient, every day.

STAGE 1

Full Schedule Intake

Every patient on the day's schedule enters automatically — no manual triage, no one excluded early — before the clinic opens.

STAGE 2

Broad Extraction Filter

Fast, rule-based matching of structured data — visual acuity ranges, prior treatment history — against protocol criteria. Casts a wide net to avoid premature exclusion.

STAGE 3

AI / LLM Deep Filter

The core differentiator. An LLM reads free-text chart notes; image models quantify pathology from OCT and retinal images. Both map against inclusion, exclusion, and study-eye criteria.

STAGE 4

Physician Flag & Discussion

Candidates who clear the deep filter surface to the treating physician with a concise eligibility summary at the point of care. The physician makes the final call.

STAGE 5

Official Screen & Randomization

Confirmed candidates proceed to formal screening already well pre-qualified — which is why screen-fail rates drop and coordinators focus on genuine candidates.

Stage 4: Physician Review

Patients who clear the deep filter arrive with a short eligibility summary delivered at the point of care. In our workflow, the candidate list also goes out the evening before clinic as a secure summary to the treating physicians and technicians, and a brief eligibility note is placed in the chart (Figure 2). The physician reads the reasoning and makes the final call. The trial conversation is prompted by evidence, and clinic flow stays intact.

Stage 5: Formal Screening

Confirmed candidates move to formal, pre-qualified screening. Using this process, screen failures drop and coordinators can spend their effort on genuine candidates instead of long shots.

THE TREATMENT-NAIVE CONSULT

Many protocols require a treatment-naive study eye, and those patients arrive as new consults with thin records and a one-visit window. A referral note may be all the chart holds the night before, and if the visit ends with a standard-of-care injection, eligibility is gone. Thus, the system runs a second, faster loop for new patients. Referral packets are parsed the evening before, and consult slots whose visit reason suggests a naive-eligible diagnosis are placed on a watch list. The screening and trial conversation then happen on the day of the visit. OCT and fundus images are quantified within minutes of capture, and the eligibility flag reaches the physician before the treatment decision is made, with a reminder that any same-day injection closes the protocol window.




Figure 2. This is an illustrative point-of-care view of the AI pre-screener, which has flagged candidates by eligibility status and model confidence for physician review. All patient names, medical record numbers, and study names shown are fictional and for demonstration only.

COMPLIANCE CONCERNS

Reviewing existing records to identify prospective participants is a well-established recruitment activity that Institutional Review Boards (IRBs) routinely approve, and HIPAA allows it under the preparatory-to-research provision. The review must stay inside the practice, distribution must be secure and limited to the care and research team, and no protocol-specific procedure can happen before informed consent is signed.

With our system, the treating physician makes first contact. The AI's output is a flag for review, never a determination, and the chart language reflects that, as patients often read their own notes. Every pre-screen is logged automatically, which sponsors increasingly expect, and patients who decline are marked so they are not approached again. Sites adopting a similar workflow should confirm that AI-assisted pre-screening is described in their IRB-approved recruitment plan and that any vendor handling patient data operates under a business associate agreement.

WHAT ACTUALLY CHANGES

In practice, systematic extraction flags more eligible candidates than manual review. Deep filtering and image quantification reduce screen failures, because much of the disqualifying information is caught before anyone is consented. In our practice, this structured onboarding process can bring a protocol live in about 4 weeks. These numbers will vary by site and protocol, and every practice should test them against its own data.

The underlying philosophy is simple: AI does the heavy lifting, and physicians make the decisions. Patients get a fair shot at trial enrollment, practices build durable enrollment

pipelines, and cleaner candidate selection and lower screen failure rates shorten timelines and improve data quality.

Rather than requiring more patients in the practice, this workflow is designed to identify the patients we already see. Your next enrollee is likely on tomorrow's schedule. Is your process built to recognize them before they walk back out the door? ■

AI disclosure: Claude Opus 4.8 (Anthropic) was used to assist with drafting the text from an author-prepared outline. All content was reviewed, verified, and revised by the authors, who assume full responsibility for the accuracy, originality, and integrity of the manuscript. The AI tool was not used for data analysis, interpretation, or drawing scientific conclusions.

SUNIL GUPTA, MD

■ Retina Specialist, Retina Specialty Institute, Pensacola, Florida
■ sgupta@retinaspecialty.com

JOSHUA KOO, MS, BS

■ Director of Innovation, Retina Specialty Institute, Pensacola, Florida
■ jkoo@retinaspecialty.com

Financial disclosure: The AI screening system described in this article was developed at Retina Specialty Institute in partnership with Coherent Health, where both authors work. It is not a commercial product, and the authors receive no revenue from it. Coherent intends to make the system available to other retina practices as a larger offering named Engram, and the authors anticipate a future financial interest in its commercialization. Readers should consider this prospective interest when evaluating the article. To learn more or Beta test the system, contact the authors.