CASE STUDY

AI-Powered Pre-Test Risk Assessment Chatbot

for Rare Disease Genetic Testing

Rare Disease Testing · AI-Powered Pre-test Assessment

Where the Opportunity Was

The lab's testing pipeline was working well. Sequencing, variant analysis, reporting, and the clinical machinery were sound. But what arrived at intake told a different story.

A portion of referrals came in without enough structured phenotype information to clearly justify the test being ordered. Not because referring clinicians weren't skilled, they were. But a pediatrician seeing a child with developmental delay and seizures is working with a broad differential and no genetics training to help them determine whether WES is the right call, a targeted panel, or whether the clinical picture warrants testing at all.

The result, upstream, was test orders that payors pushed back on. The result, inside the lab, was genetic counselors spending significant time manually translating free-text referral notes into structured HPO terms before interpretation could even begin, a data formatting task sitting in the middle of a high-skill clinical workflow.

The lab's leadership read this clearly: the solution wasn't to process denials faster. It was to create clinical clarity earlier, at the moment the clinician first sees the patient, before the order is placed, before the insurer is involved.

That clarity, built into a tool for the referring clinician, would give everyone downstream a cleaner signal to work from.

What the Lab Set Out to Solve

Getting to a definitive picture before the order is placed

For a rare disease, the stakes of an uncertain referral are high. WES runs $2,000 to $5,000 or more. If the phenotype picture isn't strong enough to support the test, the order gets denied, and the whole cycle, review, rewrite, resubmit, begins. The lab wanted a way to surface that clinical picture at the very first encounter, so the decision to test was grounded in structured evidence before it went anywhere.

Giving referring clinicians the same phenotype logic the lab uses internally

The lab's interpretation team works from HPO-coded phenotype profiles. The computational matching tools they rely on, the same logic underlying platforms like Exomiser and Phenomizer, require structured phenotype input to work well. That structured input was being assembled by genetic counselors after the fact. The opportunity was to move that structuring upstream, into the hands of the ordering clinician, at the point of referral.

Freeing counselors for the work that requires them

With intake translation removed, genetic counselors could return their attention to what actually needs them: variant interpretation, VUS resolution, and patient counseling. These are not tasks that can be systematized. Manual data formatting is.

What We Built and Why

We built an AI-powered chatbot that sits at the front of the referral workflow, accessible to non-geneticist clinicians before they place a test order. As the clinician walks through the patient's presentation, two things happen in real time.

> **The design principle: move clinical quality control to the point of referral, not the point of intake.**

Phenotype Structuring at the Point of Care

Every symptom and clinical feature the clinician describes is mapped to standardized HPO terms as the conversation unfolds. "Seizures since 8 months" becomes *Generalized onset seizures (HP:0002197)* with *Infantile onset (HP:0003593)*. "Not talking yet" becomes *Absent speech (HP:0001344)*. The clinician stays in their clinical language. The system builds the structured profile behind it.

Phenotype-Driven Risk Scoring

The assembled HPO profile is run against known gene-disease associations using computational phenotype matching, the same evidence framework that the lab's own interpretation team relies on. The output is a ranked list of candidate conditions with match scores. The ordering clinician now has a clear, evidence-based signal: whether the clinical picture supports testing, and if so, what kind.

By the time a test order is placed, the lab already has a structured phenotype profile attached to it. And by the time it reaches a payor, the clinical justification is grounded in ontology-coded evidence, not reconstructed from memory under time pressure.

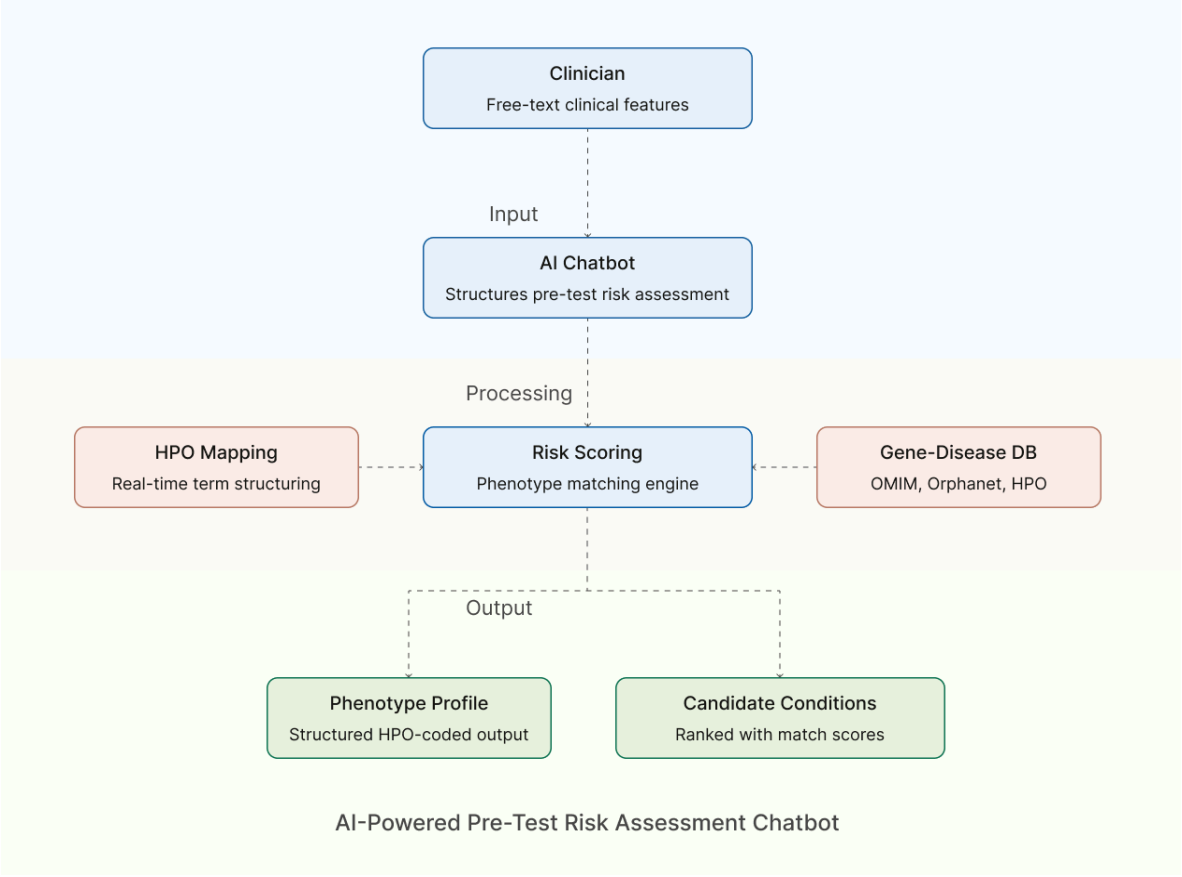

System architecture: from clinician input to lab outcomes

The clinical justification now arrives complete. Built on the same evidence frameworks, payors recognise, structurally harder to deny on medical necessity grounds.

By the Numbers

Metric	Before	After
Referral phenotype quality	Free-text, unstructured	**HPO-coded at the point of referral**
Counselor time on intake translation	2-4 Hours per day	**Significantly reduced**
Ordering accuracy	Inconsistent	**Phenotype-guided**
Ordering decision basis	Clinical intuition, free-text notes	**Phenotype-driven risk score**
Diagnostic yield	Variable	**Improved through better referral filtering**

The Impact

- **Clarity before the order, not after the denial.** The chatbot delivers a structured, evidence-based assessment at the first clinical encounter.
- **Cleaner data flowing into the lab from day one.** HPO mapping now happens during the clinical encounter. The interpretation team receives complete, structured phenotype profiles at intake rather than assembling them from unstructured notes. Variant prioritization improves as a result.
- **Counselors are back on clinical work.** Intake translation is no longer part of the genetic counselor's day.
- **Patients on the right path sooner.** When the ordering decision is grounded in structured phenotype evidence from the first encounter, patients with actionable genetic conditions reach a diagnostic result faster.

Phenotype-driven risk scoring using HPO terms is established clinical science. What NonStop built is the conversational interface that makes it available to every referring clinician, at the point of care, before a single order is placed.

The lab didn't wait for denials to tell them something was wrong upstream. They built clinical clarity into the earliest moment of the referral process, and the rest of the workflow improved because of it.

Interested in building the same capabilities for your lab?

nonstopio.com | Clinical Genomics & Life Sciences Engineering