

HHS/NIH TOPx Tech Sprint · Invisible Illness track · October 2026

Undiagnosed No More

The Invisible Patient Finder

Find the patients the health record misses: open, tested tools for patients, clinicians, health systems, and developers.

Amir Asiaee

Vanderbilt University Medical Center

Kaveh Aryan

Department of Informatics, King's College London

undiagnosednomore.org

Some illnesses are real but invisible in the record

Long COVID, ME/CFS (chronic fatigue), fibromyalgia, and Lyme disease are hard to diagnose. Even after a clinician names the illness, the diagnosis code is often missing from the chart.

Without the code, the next clinician, post-COVID clinics, benefit programs, clinical trials, and studies that count patients by code can all miss the patient.

The diagnostic odyssey


In the roughly 100 days before their Long COVID was coded, patients had

6.6 more outpatient visits

0.4 more emergency visits

than matched post-COVID patients, about **\$900 per patient** at Medicare prices, not counting missed work.

Long COVID is coded for about 1 in 9 adults who ever had it



We compared how many adults say they have had Long COVID, in a national CDC survey, with how many carry its diagnosis code in their health record.


Less often for women

Women report Long COVID more often than men, but a smaller share of their Long COVID is coded: **0.79 times** the share for men (95% interval 0.65 to 0.95). This comparison is not adjusted for health system.

Across the 10 health systems with enough patients, coding ranged from 0.79 to 2.37 times the expected rate (10th to 90th percentile).

Diagnoses patients report are often missing from their record

People told a survey that a clinician had diagnosed them. How often does their record carry the code, compared with high blood pressure and type 2 diabetes (coded as often = 1)?



Share coded among people who report the diagnosis, relative to high blood pressure and type 2 diabetes, which are coded for 83% to 90% of the people who report them. Chronic fatigue at 0.28 means it is coded less than a third as often.

Source: NIH All of Us Research Program; 114,610 survey respondents with at least 5 years of health records before the survey (161,480 respondents overall), among those reporting ongoing care for each condition. Companion preprint on coding of self-reported diagnoses (in preparation).

A summary-count model can leave the research enclave

All of Us data never leave its secure workspace, and neither may a model trained on individual records. So we fit our model only from privacy-protected summary counts, which the program allows us to export.

The result is a small, open file: a score over named, dated chart findings (new symptoms, tests, prescriptions, repeat visits) that another health system can run and check against our test patients.

0.66

AUROC at **18 health systems** held out from fitting (the summary-count score refit without them), against 0.62 for a full-record model that cannot leave.


55 days

Median lead before the Long COVID code, among later-coded patients in the top 5% of scores (full-record model).

AUROC is the chance the model ranks a patient who is later coded above one who is not: 0.5 is a coin flip, 1.0 is perfect.

Training on patients' reports ranks uncoded patients higher

Most models learn only from patients who already have a code. We also trained on patients' own reports of a diagnosis, then asked which model ranks higher the people who report the illness but have no code for it.




AUROC for ranking people who report the diagnosis but have no code. For endometriosis, discrimination was low and the report-trained score did not predict later codes (0.50 against 0.70). Women who reported it were older and had fewer recent menstrual and hormonal records, so an ever-told diagnosis may describe past illness; this was not tested.

Source: NIH All of Us Research Program, a random 20% validation split of each cohort (149,860 survey respondents for chronic fatigue, 150,340 for fibromyalgia, 97,590 women for endometriosis). Companion preprint on case-finding from patients' reports (in preparation).

How it fits together

One engine, four doors



The patient app: get ready for the visit

The screenshot shows a patient app interface for a user named Maria with a 'Long COVID pattern'. The app is titled 'Patient app' and is an 'Interactive demo · synthetic answers'. The page is 'Page 14 of 14: Summary' and includes 'Back' and 'Done' buttons. The main heading is 'Your summary', followed by the text 'These answers can guide a clinician's review.' and 'How your answers relate to conditions'. There are four summary items, each with a 'Discuss' button:

- Long COVID**
RECOVER symptom score **16 points**
This simplified score reaches the Long COVID research cutoff of 11. Review it with your clinician; it may run high without severity questions.
- ME/CFS and Long COVID**
Symptoms after activity (DePaul screen) **Screen positive**
Worsening after activity is a key ME/CFS symptom and can occur with Long COVID. Tell your clinician what triggers it.
- Long COVID follow-up**
Daily activity after COVID-19 (PCFS) **Grade 3: moderate functional limitations**
This grade describes daily limits after COVID-19, not their cause. Ask your clinician to review them.
- Long COVID**
How long symptoms have lasted **Check matched**

- Validated questions: the NIH RECOVER symptom questions, the DePaul questions on crashes after activity (post-exertional malaise), condition self-checks, PHQ-9, and GAD-7, in English or Spanish.
- A summary of what each answer means for a condition and what to discuss. Patients never see a probability.
- A visit summary to print, or to download as a FHIR file that clinic software can read.
- Answers stay on the device. No account, no tracking.

The clinician view: the evidence inside the chart

The screenshot displays the UNOM Clinician app interface. At the top, the navigation bar includes the UNOM logo and links for Undiagnosed No More, Clinicians (active), Patients, Policymakers, Developers, and About. A help icon and settings icon are on the right. Below the navigation bar, a text block explains the app's purpose: "Inside the health record, see which condition a chart pattern points to, the dated findings behind it and how much each one moves the estimate, the patient's own answers combined into it, and what to check next, for all nine conditions." This is followed by the heading "Explore the clinician app" and a link to "Launch via SMART on FHIR".

Below this, there are four buttons for different patient scenarios: "Maria COVID-19, then new symptoms", "Elena Years of fatigue and dizziness", "Sofia Years of widespread pain", and "Dana Knee pain, facial weakness". The "Maria" button is selected.


The main content area shows the "Clinician app" interface for "Maria Synthetic", a 44-year-old female as of June 1, 2024. It features a "Needs review" section with a "Long COVID High" finding, noted as "15.4x average". The text below this finding states: "Top findings: 3+ symptom types coded in days 31-90; Personal history of COVID-19 coded (days 31-90)". Below this, a section titled "Also checked" lists other conditions: ME/CFS typical, Fibromyalgia typical, Endometriosis typical, Lyme typical, and others.

To the right, a "From the patient" section indicates "Completed synthetic responses. Answered May 28, 2024." with a "Remove" link and a "Review responses" button.

At the bottom, a horizontal tab bar allows switching between different conditions: Long COVID (selected), ME/CFS, Fibromyalgia, Endometriosis, Lyme, Lupus, Migraine, Depression, Anxiety, and Timeline. The "Long COVID" tab is active, showing a "High: chart and answers well above the average" finding. The text below this finding states: "Chance of a first Long COVID code (U09.9) 91 to 455 days after COVID-19, compared with the average post-COVID patient at day 90: 15.4 times the average (estimated 15.9%) with the patient's answers; chart alone 6.1 times (6.2%)."

- One card per condition, nine in all: Long COVID, ME/CFS, Lyme disease, lupus, fibromyalgia, migraine, endometriosis, depression, and anxiety.
- Every card lists the dated chart findings behind it and the next step, such as a post-COVID clinic referral.
- Built to open inside the EHR through SMART on FHIR, with no new login; tested in the public SMART sandbox.
- Research decision support. It surfaces patterns for clinician review.

The Atlas: who goes uncoded, and what finding them costs



Clinician time per later-coded patient found

\$116

Two-step screen: a short questionnaire first, then chart review

\$233

Chart review of every flagged patient

The calculator assumes each patient found is recognized about 55 days before the code.

Atlas calculator defaults: 10,000 post-COVID patients screened, the top 5% flagged (measured recall 19% and PPV 4.3%), 5 minutes per review at \$120 an hour, 40% pass the questionnaire. Change any of them in the Atlas.

Open standards, and health data that stay local

Standards

- **FHIR R4** for reading the chart
- **SMART on FHIR** to open inside the EHR
- **CDS Hooks** to send a card when a chart opens
- Standard codes: ICD-10-CM, LOINC, RxNorm, CPT

Where health data stay

- The clinician view scores the chart in the browser.
- The CDS Hooks service code keeps no copy of the chart data it receives and does not log it.
- The patient app keeps answers on the device.
- Every demo patient is synthetic.

All of Us disclosure rules

- Only checked, aggregate numbers leave the research workspace.
- Counts of 1 to 20 are suppressed, and other counts are rounded up to a multiple of 10.
- The models were fit from those summary counts, so no individual record left.

Open model files, a reference scorer, and test patients let another health system check that its copy reproduces our scores.

Roadmap and ask

Works today

Today

Patient app, clinician view, Atlas, open models, and CDS Hooks, with synthetic patients. Launches from the public SMART sandbox; registered with Epic and athenaOne, with sandbox access pending.

Next

Pilot

A silent-mode pilot at a health system (cards scored but not shown) to check them against real charts, and validation in a second national research network.

Planned

Marketplace

Listings in EHR app marketplaces, and new conditions on the same engine, such as POTS and post-treatment Lyme disease syndrome.

Our ask: help us move from sandbox to a health system. A pilot site, EHR marketplace review, and patient advocates to review every new card.