EVIDENTLY

The 4.5 Missing Comorbidities: How Clinical Data Intelligence is Rewriting the Admit Note

Every admit note is a bet that the provider found everything that matters in the chart. At most health systems, that bet is losing an average of 4–5 quality-impacting conditions per patient.

Researchers at the University of Iowa Health Care published the first peer-reviewed study of AI-generated admission notes in a real-world inpatient setting, using Evidently's Clinical Data Intelligence platform. The finding: a median of 4.5 quality-impacting conditions per chart — present in the longitudinal record but absent from the provider's admit note[1]. Conditions buried in scanned documents, outside records, health information exchange data, and historical notes the EHR was never built to understand. At 97% accuracy.

We'll examine that study alongside cross-site validation at four other leading health systems — different geographies, different documentation cultures, one consistent, repeatable outcome.

4.5	97%	5
NEW QUALITY-IMPACTING CONDITIONS PER CHART	ACCURACY ON NET-NEW DOCUMENTED CONDITIONS	INDEPENDENT HEALTH SYSTEMS WITH CONSISTENT, REPEATABLE FINDINGS

---

[1] Rodrigues AC, Misurac J, Knabe LA, Barker K, Blum JM. Generative AI in Admission Notes and Diagnostic Completeness: A Pilot Study. *Appl Clin Inform.* 2026 May 18. DOI: 10.1055/a-2876-0998. PMID: 42150573.

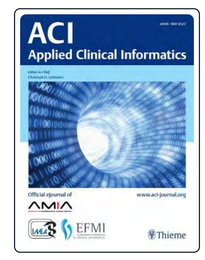

1. If You Can't Read the Whole Record, You Can't Document Patient Complexity Accurately.

Every inpatient leader knows the problem. A massive amount of patient complexity is buried in the chart — in outside records, scanned documents, faxed referral packages, Care Everywhere feeds, and imaging reports stored in systems like OnBase. About 30–40% of the chart lives in formats the EHR was never built to understand. For the provider writing the admit note, these records might as well not exist.

The conditions trapped in that 30–40% are not minor. They're the chronic comorbidities, the historical complications, the severity indicators that drive mortality benchmarking, Case Mix Index, and reimbursement. They're the difference between an admit note that reflects the patient's true complexity and one that reflects only what the provider had time to find.

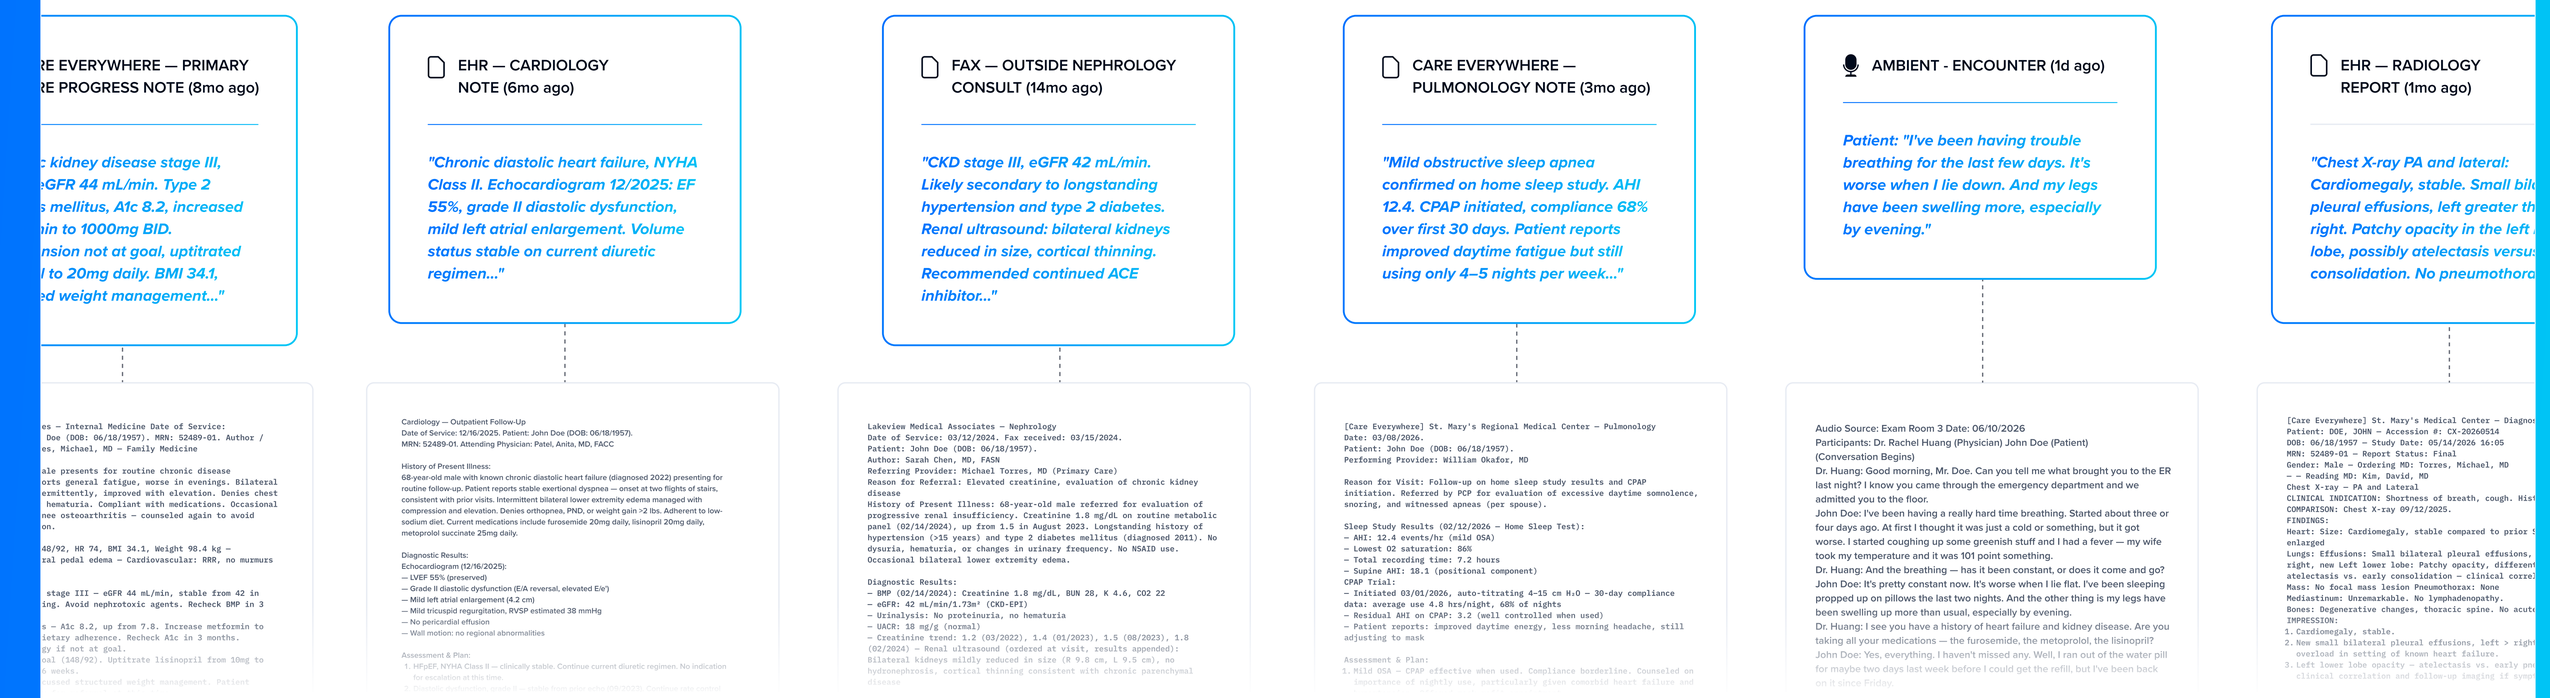

Health systems try to close this gap every day — pulling every lever available, even when the lever wasn't built for the job. Dot phrases and note templates get repurposed as behavioral nudges, embedding reminders of what to document. Pre-rounding burns hours every day on chart biopsy. Query templates chase what the admit note missed — after the note is already signed. But you can't nudge, pre-round, or query your way to complete documentation when the longitudinal complexity is buried so far in the chart that no one can find it.

None of these approaches address the root cause. The documentation gap at admission isn't a behavior problem, a technology gap in encounter capture, or a timing problem in the coding workflow. It's an architectural problem: the most consequential patient data lives in the hardest-to-reach part of the record.

2. The University of Iowa Health Care Study

In 2026, researchers at the University of Iowa Health Care set out to measure this gap — and what happens when AI reads the full record before the provider writes.

STUDY DESIGN

The team retrospectively reviewed matched pairs of admission notes across ICUs, general medicine, and intermediate care units, focusing on Assessment & Plan contents — the section of the admit note that drives severity adjustment, coding, and quality benchmarking. For each admission, Evidently produced an AI-drafted admit note approximately three hours post-admission using only pre-existing chart data — orders, labs, imaging, prior notes, ED documentation, and health information exchange data. The provider's own note was explicitly excluded from Evidently's input to prevent contamination. Every net-new documented condition surfaced by Evidently was adjudicated for clinical validity by the institution's Clinical Documentation Integrity team. The University of Iowa IRB reviewed the study and classified it as quality improvement.

KEY FINDINGS

> **Median 4.5 net-new quality-impacting conditions per chart**
>
> (IQR 2.3–6), totaling 94 across the 22 charts.[1]

> **95% of charts had undocumented conditions**
>
> 21 of 22 admissions contained at least one quality-impacting condition present in the Evidently-generated note but absent from the provider note.[1]

> **97% accuracy on net-new documented conditions**
>
> Of 104 net-new secondary documented conditions, 101 were adjudicated as clinically valid. No hallucinations were found to be present.[1]

---

[1] Rodrigues AC, Misurac J, Knabe LA, Barker K, Blum JM. Generative AI in Admission Notes and Diagnostic Completeness: A Pilot Study. *Appl Clin Inform.* 2026 May 18. DOI: 10.1055/a-2876-0998. PMID: 42150573.

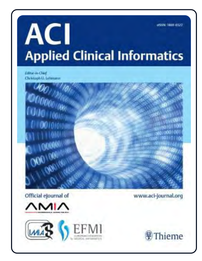

THE CASE FOR CLINICAL DATA INTELLIGENCE AT UI HEALTH CARE

These conditions were missed because they're trapped in parts of the record that are nearly impossible to reach — in outside documents, scanned faxes, Care Everywhere feeds, imaging reports, and historical data stored in systems the EHR was never built to understand.[1] Capturing them requires reading the entire patient history across every source, including the 30–40% of the chart that arrives in formats the EHR was never built to understand. Mortality benchmark-impacting factors were the largest category surfaced — 30 of 94 — the inputs that directly drive quality rankings.

For Clinical Documentation Integrity teams, the implication is a structural workflow shift. Most operate retrospectively, querying after the note is signed. Evidently surfaces these documented conditions at the point of admission — before Clinical Documentation Integrity intervention — shifting the model from retrospective remediation to concurrent capture. And empowering documentation teams to focus on high-impact queries.

UNIVERSITY OF IOWA HEALTH CARE STUDY: IMPACT CATEGORIES

CATEGORY	COUNT
Vizient Mortality Factors	30
Complications or Comorbidities (CC)	22
Elixhauser Comorbidities	19
Hierarchical Condition Categories (HCC)	13
Comorbidity / Mortality Flags	8
Major Complications or Comorbidities (MCC)	2

[1] Rodrigues AC, Misurac J, Knabe LA, Barker K, Blum JM. Generative AI in Admission Notes and Diagnostic Completeness: A Pilot Study. *Appl Clin Inform.* 2026 May 18. DOI: 10.1055/a-2876-0998. PMID: 42150573.

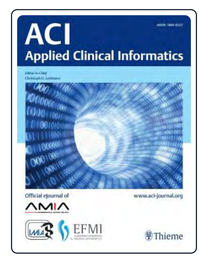

3. Cross-Site Validation at 4 Health Systems

Separately, Evidently has studied this same impact at 4 additional health systems. At each site, Evidently produced AI-drafted admit notes using only pre-existing chart data — then compared them against provider-authored notes for the same patients. Every net-new documented condition was adjudicated for clinical validity. Four health systems, four independent studies — different geographies, different documentation cultures, one consistent, repeatable outcome.

Site A: Major IDN

- Regional network
- Broad payer mix
- Community and tertiary hospitals
- Predominantly community-based care across urban and rural facilities

+24%
MCC ACCURACY

+60%
CMS MORTALITY

Site B: Large Academic Medical Center

- Tertiary referral
- Heavy inflow of outside records
- Multi-facility system spanning academic and community hospitals

+41%
MCC ACCURACY

+79%
CMS MORTALITY

Site C: Regional Health System

- Multi-state system
- Mixed-EHR footprint
- High transitions volume

+25%
MCC ACCURACY

+46%
CMS MORTALITY

Site D: Major Academic Medical System

- Multi-hospital system
- Urban academic and community sites
- Dense, high-acuity patient population

+26%
MCC ACCURACY

+52%
CMS MORTALITY

4. Unlocking Longitudinal Complexity

The documentation gap at admission is a symptom. The condition is broken access to the longitudinal patient record. Today, every tool in a health system's technology stack has built its own incomplete view of patient complexity — none of them complete, none of them consistent. For clinicians, this means as many versions of a patient's true history as there are tools in their EHR ecosystem.

ENRICHING THE WORKFLOWS CLINICIANS RELY ON

Clinical Data Intelligence addresses this at the root cause level: index the entire longitudinal patient history — including the scanned documents, outside records, faxed referral packages, and HIE feeds other tools can't read — map it through a comprehensive medical knowledge graph, and surface that intelligence into the workflows clinicians and revenue teams already use. One indexed record. One source of clinical truth. Every workflow reads from the same foundation.

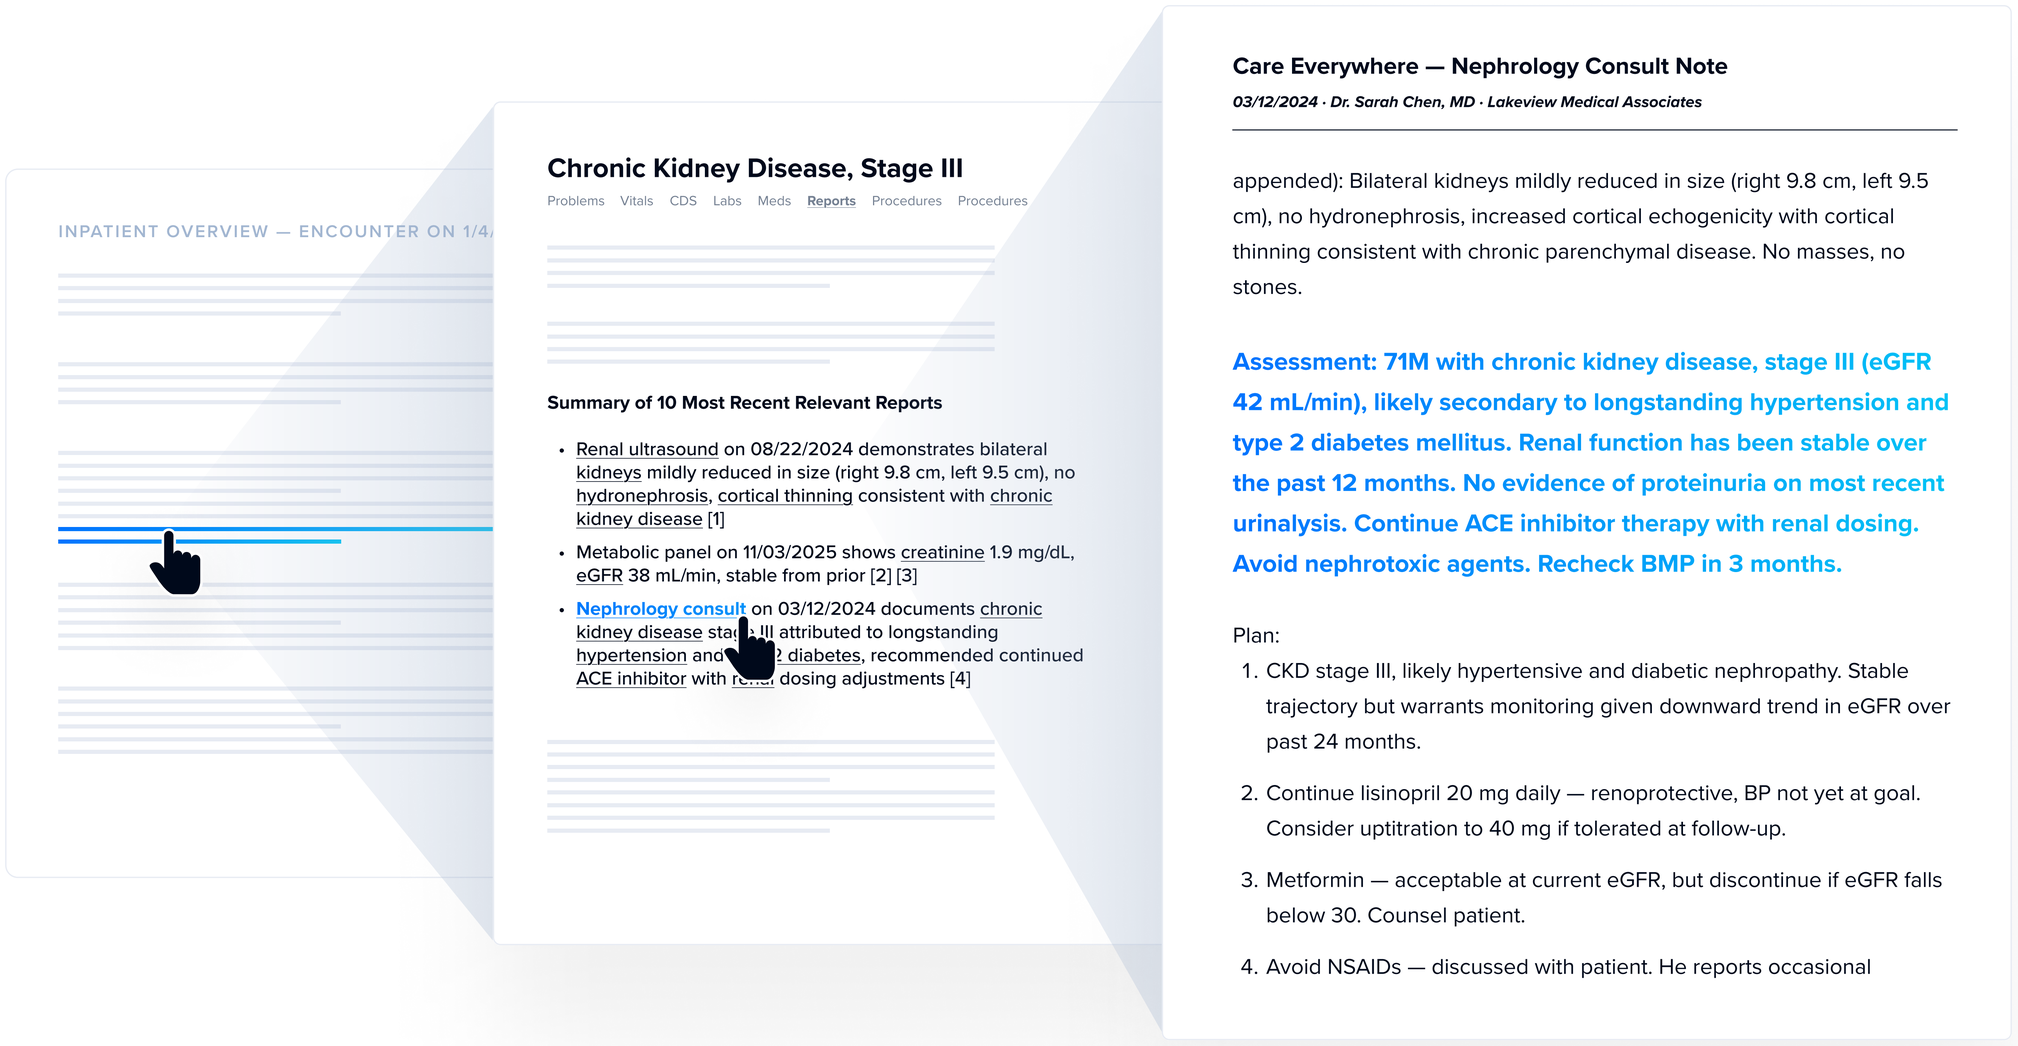

REPEATABLE RESULTS, REAL IMPACT

The evidence presented here demonstrates what happens when that foundation is applied to inpatient admission documentation — repeatable impact across five independent institutions. But the same architecture powers workflows across the health system. Understanding a patient's history better than previously possible drives repeatable outcomes.

TOOLS FOR EVERY STAGE OF THE INPATIENT DOCUMENTATION LIFECYCLE

The same Clinical Data Intelligence foundation that powers these admit note results serves the entire inpatient documentation lifecycle — from the first note to the final appeal.

Supercharging Providers by Surfacing Longitudinal Complexity at the Note

AI Drafting the Admit Note

Full patient complexity captured right when the documentation starts. An average of 4–5 quality-impacting conditions per chart, and increases in CMS mortality accuracy of up to 79%.

AI Drafting the Progress Note

The full record re-read daily — not carried forward from admission. New outside records, labs, and imaging reflected in each day's note. No documentation drift, no conditions lost mid-stay.

AI Drafting the Discharge Summary

The complete clinical picture, recapped and source-traceable. CMI reflects true patient complexity, ready for coding and quality review the moment the patient leaves.

AI-Powered Clinical Documentation Integrity Review

A suite of documentation integrity tools embedded directly in the EHR, and tightly coupled with access to full longitudinal complexity for finding the most challenging conditions and surfacing them as high-yield queries

AI-Drafted Denial Appeals

Every condition source-traceable with patent-pending concept hyperlinking. Appeals start with the evidence, not a chart biopsy.

Custom Summaries & Ask Evidently

Any clinician, any question, anywhere in the workflow. It’s like a great resident that’s read everything in the chart and is ready for any question you can throw at them.

> *My single biggest regret is not turning it on sooner.*

JAMES BLUM, MD CDH-E

CHIO @ UNIVERSITY OF IOWA HEALTH CARE

See what your admit notes are missing.

Take the H&P Compare Study.

The H&P Compare Study is the operational version of the evidence presented here: a one-week prospective study that runs in the background with zero workflow change. The output is a written clinical study with chart-level findings, benchmarked against results at leading health systems.

One week. Zero lift. Concrete findings your leadership team can act on.

hello@evidently.com