Featherglass
HEALTH

CLINICAL INTELLIGENCE BRIEF · APRIL 2026

Beyond the Scribe: Why First-Generation AI Tools Don't Fully Fix Behavioral Health — and What Comes Next

Executive Summary

Ambient AI scribes are a genuine step forward. The largest peer-reviewed study of their kind, published in *JAMA* in April 2026, confirmed that AI scribe adoption is associated with meaningful reductions in documentation time — roughly 16 minutes per 8-hour shift — and modest increases in visit volume. These are real gains, and organizations that haven't explored ambient documentation tools should.

But for behavioral health organizations operating across levels of care, the data also reveal a harder truth: documentation time is not the primary driver of clinician burnout, administrative overload, or unsafe care transitions. The larger disease — the one that 1st-generation scribes were never designed to treat — is the fragmentation of the patient record and the systemic amnesia it produces.

Fragmentation is the structural reality that patient records are created, stored, and siloed separately across every facility, EHR, and level of care a patient encounters. Systemic amnesia is its clinical consequence: the behavioral health system behaves as if it has no memory of a patient's prior care, forcing every new encounter to begin from scratch. The 2026 *JAMA* study found that clinicians who adopted AI scribes reallocated much of their saved documentation time to other administrative tasks — including "reviewing current or prior documentation for accuracy... conducting medical record review." In other words, the time savings from transcription were largely absorbed by the very problem transcription cannot solve.

This brief explains why, for behavioral health, solving documentation is necessary but not sufficient — and what a complete clinical memory infrastructure actually requires.

01. The Documentation Crisis Is Real

Before examining the limits of 1st-generation AI scribes, it is important to be clear about what they do well. The EHR documentation burden on clinicians is well-documented, substantial, and harmful. Clinicians average 2.3 hours of EHR time per 8 hours of patient care. After-hours EHR activity — so-called "pajama time" — is independently associated with burnout. Reducing the time a clinician spends dictating and editing notes has real value: it reduces cognitive load, may increase presence during patient encounters, and can expand clinical capacity.

Ambient AI scribes — tools that listen passively during encounters and generate structured notes from natural clinical dialogue — address this problem directly. They do not require a change in clinical workflow. They do not require manual dictation. They produce notes in real time, and they have been shown in multiple studies to improve clinician satisfaction and reduce the sense of documentation burden.

For organizations that have not yet adopted any form of AI-assisted documentation, ambient scribe technology is a meaningful and worthwhile investment. The documentation problem it addresses is real.

The question is what happens next — and why, for behavioral health organizations in particular, the gains from transcription are bounded by a deeper structural problem that ambient scribes are not built to solve.

02. What the Largest Study Ever Published Actually Found

In April 2026, *JAMA* published the most comprehensive evaluation of AI scribe adoption to date. Rotenstein and colleagues studied 8,581 clinicians — 1,809 AI scribe adopters and 6,772 non-adopters — across five US academic health systems over a period spanning June 2023 to August 2025. The study used a difference-in-differences design to isolate the association between AI scribe adoption and changes in EHR time, documentation time, after-hours work, and visit volume.

What the Study Confirmed

AI scribe adoption was associated with meaningful and statistically significant findings. For primary care clinicians — those with the highest documentation burden — the documentation time reduction was even larger, at 26.9 minutes. Clinicians who used AI scribes in more than half of their encounters saved 27.3 minutes of documentation time, suggesting a dose-response effect.

16.0

fewer minutes of documentation time per 8 hours (a 10% relative reduction)

Rotenstein et al., JAMA 2026

13.4

fewer minutes of total EHR time per 8 hours (a 3% relative reduction)

Rotenstein et al., JAMA 2026

0.49

additional weekly visits delivered by adopting clinicians

Rotenstein et al., JAMA 2026

$167

estimated additional monthly EM revenue per adopting clinician

Rotenstein et al., JAMA 2026

What the Study Did Not Find

The study's most important finding for behavioral health leaders may be what *did not* change.

No Significant Change

After-hours "pajama time" — work outside scheduled hours — did not change significantly. The adjusted difference was 3.1 minutes, with a confidence interval that included zero.

Rotenstein et al., JAMA 2026

This is not a rounding error or a study limitation; the result was robust across multiple sensitivity analyses and replicated the findings of prior single-site studies. Clinicians who adopted AI scribes went home and opened their laptops at the same rate as those who did not.

Where Did the Time Go?

The authors offer a direct explanation: clinicians who saved time on documentation appear to have reallocated much of that time to "other patient care activities, such as reviewing current or prior documentation for accuracy, answering electronic inbox messages from patients, addressing test results, or **conducting medical record review.**"

> "The question is no longer whether ambient AI can save documentation time for clinicians but whether that time is reinvested in ways that measurably improve outcomes."
>
> — Tierney, Lee, and Liu, JAMA Editorial (April 2026)

Put differently: transcription reduced one administrative burden. Much of that freed capacity was immediately absorbed by another — the burden of manually reconstructing the clinical narrative from fragmented, siloed records. The pipeline got faster at one end; the bottleneck simply moved downstream.

03. Two Problems Transcription Cannot Solve

The 2026 *JAMA* findings are not a condemnation of ambient scribes. They are a precise map of what those tools do and do not address. To understand the gap, it is necessary to name the two underlying problems that transcription is not designed to touch: **fragmentation** and **systemic amnesia**.

Fragmentation: The Structural Problem

Fragmentation describes the reality that patient records are created, stored, and siloed separately across every care encounter, institution, EHR system, and level of care. A patient who has received inpatient psychiatric care, followed by residential treatment, followed by partial hospitalization at a separate facility, has a clinical record distributed across at least three distinct systems — often in incompatible formats, behind different access credentials, and practically inaccessible to the clinician conducting the next intake.

Fragmentation is not a failure of documentation effort. Clinicians and care teams often document thoroughly. The problem is that thorough documentation in System A does not travel to System B when the patient transitions. The record exists. It is simply unreachable when it is most needed.

Systemic Amnesia: The Clinical Consequence

If fragmentation is the structural problem, systemic amnesia is what that structure produces at the point of care. Systemic amnesia describes the condition in which the behavioral health system — despite years of prior engagement with a patient — behaves as though it has no memory of that patient's clinical history.

The result is a care system that perpetually begins from scratch. Each new admissions clinician reconstructs a history that has already been reconstructed at three prior facilities. Each new treatment team re-establishes diagnostic impressions that prior teams reached, documented, and then lost. Each level-of-care transition carries the clinical narrative only as far as the receiving fax machine — which is to say, not far enough.

Systemic amnesia is not a technology failure in the narrow sense. It is the cumulative effect of thousands of individual handoffs that had no shared memory layer connecting them. The patient's story exists in fragments. No one has ever assembled it into a continuous, portable, living record — because no infrastructure existed to do so.

Why These Problems Are Compounding

Fragmentation and systemic amnesia are distinct but self-reinforcing. Fragmentation creates the conditions for systemic amnesia; systemic amnesia ensures that fragmented records are never synthesized into something clinically usable. The behavioral health system thus develops a structural bias toward forgetting — not because clinicians are careless, but because the infrastructure was never designed to remember.

An ambient AI scribe is exceptionally good at one thing: capturing what happened in the room during this encounter. It does not know — and was never designed to know — what happened in the room during the last 12 encounters, at three other facilities, with four different treatment teams.

04. Why Behavioral Health Is Uniquely Vulnerable

Every clinical specialty struggles with fragmentation to some degree. Behavioral health faces it with unusual severity for several interrelated reasons.

Clinical Narratives Are Longitudinal by Nature

In behavioral health, the clinical record is not a snapshot — it is a timeline. A psychiatrist's understanding of a patient's presentation depends heavily on the arc of prior episodes: what worked, what didn't, what medications caused adverse effects, what trauma disclosures have already been made, what the patient's family reported during a prior crisis. This longitudinal context is not supplementary to care quality. It is the foundation of it.

An ambient scribe captures what happens in the current encounter with precision. But a current encounter with no access to prior narrative is clinically impoverished in a way that does not apply to the same degree in, say, an orthopedic visit.

The Step-Down Continuum Creates Acute Transition Points

High-acuity behavioral health care is structured around a continuum of levels: inpatient stabilization, residential treatment, partial hospitalization (PHP), intensive outpatient (IOP), and outpatient maintenance. Each transition down this continuum is a potential moment of catastrophic information loss.

When a patient steps down from a residential program to a PHP, the receiving clinical team typically has access to a discharge summary and, if they are fortunate, a follow-up call from the discharging clinician. What they do not have is a synthesized, searchable, annotatable record of 30 to 90 days of intensive treatment — the session notes, the treatment team discussions, the medication trials, the family contact records, and the patient's own documented progress.

The admissions team at the receiving program rebuilds this history manually. This process takes hours. It is error-prone. It is a direct drain on clinical capacity. And it is repeated at every transition, for every patient, indefinitely.

The Harm of Retelling

For patients with complex trauma histories — a population that constitutes a substantial proportion of high-acuity behavioral health caseloads — the systemic amnesia problem carries a direct clinical harm that transcends administrative inefficiency. Patients with PTSD, dissociative disorders, and histories of repeated trauma are asked, at each new intake, to recount the events and experiences that constitute their clinical history. This retelling is not neutral. It can be retraumatizing. It can rupture therapeutic alliance before it has had the chance to form.

The argument for a longitudinal, portable clinical record is not simply operational. For this population, it is an ethical one.

The Admissions Reconstruction Tax

For organizations operating at scale — multi-site networks, systems coordinating across residential and outpatient facilities, programs receiving regular step-down referrals — the cumulative cost of systemic amnesia is staggering. Consider the arithmetic: if an admissions team spends two hours per new admission reconstructing clinical history from fragmented prior records, and a facility processes 200 admissions per year, that is 400 clinician-hours annually dedicated to work that should not be necessary. At scale across a network, the number compounds rapidly. This is clinical capacity that is being consumed not by care, but by the absence of memory infrastructure.

05. What a Complete Solution Requires

None of the foregoing is an argument against ambient documentation tools. It is an argument for thinking clearly about what those tools address and what they do not — so that organizations can build toward a complete solution rather than a partial one.

A complete solution to behavioral health's documentation and continuity crisis requires three distinct capabilities, each building on the last.

Capability 1: Ambient Encounter Capture

The first capability is what ambient AI scribes already provide: capturing natural clinical dialogue and generating structured, compliant documentation from it — progress notes, intake evaluations, daily assessments. This is the foundation. Clinicians who are less burdened by real-time documentation are more present, more productive, and more likely to engage deeply with the other capabilities that follow. Every organization should have this layer in place.

Capability 2: Longitudinal Clinical Memory

The second capability is the layer that 1st-generation scribes do not address: building a continuous, evolving clinical memory from every captured encounter. Rather than producing a note that lives as an isolated document, a longitudinal memory layer connects each encounter to the ones that preceded it. Patient histories, symptom trajectories, diagnostic progressions, and prior treatment decisions remain continuous and queryable — not as a pile of documents, but as a structured clinical timeline.

This is the capability that prevents the admissions team from starting over. When a patient steps down from residential to PHP, a longitudinal clinical memory carries the full narrative of the prior

episode — synthesized, organized, and immediately accessible — to the receiving team. The patient's story doesn't disappear at the discharge date. It travels.

Capability 3: Advanced Clinical Synthesis and Cross-Institutional Infrastructure

The third capability operates at the enterprise level: enabling secure, HIPAA-compliant record sharing and clinical synthesis across different facilities, care teams, and institutions. This is not document transfer — it is a shared clinical workspace in which providers across the continuum of care can access, annotate, and build on the same longitudinal narrative.

This capability addresses the hardest version of the fragmentation problem: not the fragmentation within a single organization's records, but the fragmentation that exists between organizations. When a patient steps down from a residential program operated by one organization to an IOP operated by another, the standard of care is still a fax and a discharge summary. A cross-institutional clinical workspace replaces that with a living clinical record that travels securely — and that both the sending and receiving team can access and contribute to.

At this layer, advanced document generation becomes possible in a qualitatively different sense than what ambient scribes provide. Discharge summaries that synthesize months of treatment, step-down care narratives that integrate prior records and recent encounters, psychological testing reports that draw on a patient's complete longitudinal history — these are documents that require not transcription, but clinical intelligence. They require a system that has genuinely learned the patient's story.

06. A Framework for Evaluating Clinical Memory Infrastructure

For organizations evaluating AI-enabled clinical tools — whether considering a first investment or reassessing an existing one — the following questions can help distinguish tools that address the documentation problem from platforms that address the memory and continuity problem.

On ambient capture:

- Does the platform generate compliant, structured documentation from natural clinical dialogue across all care contexts in your organization?
- How does it handle behavioral health-specific documentation formats?

On longitudinal memory:

- Does each generated note exist as a standalone document, or does it contribute to a continuously evolving patient record?
- If a patient returns after six months, can the platform surface a synthesized summary of prior treatment?
- If a care team member changes, is clinical continuity preserved without manual handoff?

On cross-institutional infrastructure:

- Can the platform securely share a patient's clinical record with a receiving facility that uses a different EHR?
- Does record sharing mean sending a document — or providing access to a shared, annotatable clinical workspace?

On advanced synthesis:

- Can the platform generate documents that require clinical synthesis — discharge summaries, step-down care narratives — not just documentation of current encounters?
- How does it handle patients with complex, multi-episode histories?

On clinical continuity across the step-down spectrum:

- Has the platform been designed specifically for high-acuity and step-down behavioral health environments?

- Does the pricing model align with clinical volume across a care network?

Conclusion

The 2026 *JAMA* data confirm what many behavioral health clinicians have intuited for years: reducing documentation time is valuable, and ambient AI scribes deliver genuine gains. But those gains are partial. The time saved from transcription is, in substantial part, being reallocated to the manual reconstruction of patient histories from fragmented records — the very work that a clinical memory infrastructure would eliminate.

Behavioral health's documentation crisis has always had two dimensions. The first — the burden of real-time note generation — is being addressed. The second — the fragmentation and systemic amnesia that make every clinical transition a reconstruction project — is not.

The organizations that build for both will deliver better care. Their clinicians will spend more time with patients and less time chasing faxes. Their admissions teams will receive complete clinical narratives rather than reconstructed approximations. And their patients — particularly those with complex trauma histories — will be asked to tell their story once, because the system will finally be capable of remembering it.

References

Rotenstein LS, Holmgren AJ, Thombley R, et al. Changes in clinician time expenditure and visit quantity with adoption of artificial intelligence-powered scribes: a multisite study. *JAMA*. Published online April 1, 2026. doi: 10.1001/jama.2026.2253

Tierney AA, Lee K, Liu VX. Ambient AI scribes and the quintuple aim: what is counted — and what matters. *JAMA*. Published online April 1, 2026. doi:10.1001/jama.2026.3529

Featherglass

HEALTH

Schedule a clinical demonstration

featherglass.com

partnerships@featherglass.com

Built for high-acuity behavioral health facilities and the complex care continuums they navigate.

Baltimore, Maryland