APA's Committee on Mental Health IT developed a list of some questions that it believes relevant to gather information that may be useful in assessing an AI scribe tool. In July, 2025, the identified questions were sent to 31AI scribe tool developers and the answers received from those entities have been published at psychiatry.org/AIScribeTools. Note that not all entities which received the inquiry provided a response. The page will be updated as more responses are received.

These sample questions are intended to provide suggestions of information to consider when evaluating an AI scribe tool for potential use in a psychiatric practice. The sample answer responses reflect the information provided by the tool developer on a specific date; they have not been independently verified as accurate by APA and do not necessarily reflect the most current information for the tool. APA encourages each individual to conduct their own assessment before using any tool. APA does not endorse and has not investigated the quality of any particular AI scribe tool mentioned and assumes no responsibility for any result or outcome related to the use of the information or any tool.

Heidi Health www.heidihealth.com
American Psychiatric Association (APA) Question Responses - Heidi

1. How does your AI scribe handle clinical documentation in behavioral health, including progress notes, therapy sessions, and psychiatric evaluations?

- Heidi ambiently transcribes clinical encounters via a variety of modals, including a web-app, desktop app, mobile app or Chrome Extension. Following transcription, users can generate notes and documents of any structure & style, allowing for countless use-cases. Depending on their chosen template, Heidi uses the transcript along with any other context provided in the additional 'context' tab, to generate any document or note. Ranging from session notes to progress notes, psychiatric evaluations to psychological assessments (e.g. GAD-7, PHQ-9 etc). Users have full control over what the output looks like ~ structurally, stylistically and even dynamic behavior of a given document relative to the visit.

2. What measures are in place to ensure patient privacy and compliance with HIPAA and other relevant regulations when using your AI scribing tool?

- We achieve HIPAA compliance by closely aligning our data & privacy governance with HIPAA requirements and operate Heidi in a HIPAA-compliant infrastructure, leveraging a compliance platform that actively monitors and maintains our HIPAA controls in real-time. In addition, we have BAAs (Business Associate Agreements) in place with all subprocessors, ensuring that all data processing meets HIPAA requirements.

- To further protect patient information, we implement de-identification and pseudonymization techniques, removing or replacing PII/PHI where appropriate before processing. Additionally,

role-based access controls and detailed audit logs restrict and monitor access to sensitive data. We also employ end-to-end encryption (AES-256 for data at rest and TLS 1.2/1.3 for data in transit) to ensure the security of all stored and transmitted data.

3. How does your AI model differentiate between structured medical terminology and the nuances of psychotherapy conversations, such as patient emotions, nonverbal cues, and contextual meaning?

- Heidi primarily uses two categories of models:

    1. Speech-to-Text (STT) models for transcription of sessions (i.e., consults/visits)

    2. Large Language Models (LLMs) for generation of clinical notes and documents based on session transcripts

- Heidi's medically-trained STT model's word recall rate (WRR) is 91% and the median word error rate (WER) is 11%. While the accuracy of LLMs is harder to quantify due to the subjective nature of clinical documentation, the reportedhallucination rate is less than 1% of sessions based on user feedback (including CSAT scores) and based on QA/QC reports by Heidi's internal team of clinicians.

- The STT models are additionally finetuned using a vocabulary of medical terminologies, medication names, abbreviations, acronyms and eponyms including those relevant to behavioural health (not including patient or customer data). Further, the LLMs used by Heidi are extensively finetuned by our Medical Knowledge team with data prepared in-house, using prompt engineering to ensure that they perform contextual summarization of the transcript. This ensures that all LLM outputs are specialty-specific, contextually accurate and adherent to the template defined by users.

- In behavioural health contexts, where conversations are more subjective and emotionally complex, Heidi captures spoken emotional content verbatim and is capable of documenting them - for example, if a patient says "I just feel empty all the time", this phrasing is preserved in the transcript and, where appropriate, summarized, for example in sections like "Patient Mood and Affect" under "Mental State Examination" part of a Psychiatrist's template.

- Currently, Heidi does not have the ability to capture tone of voice and non-verbal cues such as facial expressions or fidgety behaviour, unless such cues are documented manually by the clinician as context within Heidi, or transcribed by the clinician.

4. What level of human oversight is required for accuracy, and how does the AI handle potential errors, biases, or misinterpretations in behavioral health documentation?

- Heidi is designed to support clinicians in generating accurate and contextually appropriate behavioral health documentation. As such, human oversight is required and all AI-generated outputs can be reviewed and edited by the clinician before being approved for use, i.e., for any clinical or administrative purpose. This ensures that practitioners maintain full control over the

documentation and that they remain the final decision-makers regarding the content of the generated notes or documents.

- Heidi incorporates several safeguards to address the risks of errors, biases, or content misalignment/misinterpretation, especially in sensitive domains like behavioral health:

    - **Structured Prompting and Template Use:** Documentation is generated based on clinician-defined templates and structured prompts, limiting the AI's scope and reducing the risk of hallucination or free-form interpretation. Using prompt-engineering techniques, rigid guardrails have been defined to ensure that outputs from the model are consistently accurate and relevant to healthcare context (including behavioural health). Within templates, clinicians can include any number of custom instructions in natural language (free-text) that allow them to control the behaviour of the model (for example, users can add instructions to ensure model makes certain interpretations only using a certain logic defined by them, etc).

    - **No Clinical Decision-Making:** By design, Heidi does not suggest diagnoses or management plans and it does not make inferences. Thisensures that any sensitive content (e.g., suicide risk, trauma history, substance use) is only documented if mentioned explicitly in the session or provided as context by the clinician.

    - **Bias Mitigation Testing:** Heidi is audited using a model surveillance system that assesses model performance using a diverse set of 320 clinical test cases that represent diverse clinical scenarios (including behavioral health). These test cases are designed to detect biases, omissions, or inconsistent handling of sensitive topics. Apart from objective metrics such as hallucination rates, results are also subjectively reviewed by our medical knowledge (MK) team.

    - **Feedback and Corrections:** If an error or bias is detected in the output, clinicians can flag the session for review. These cases are triaged by the MK team adjustments are made through prompt engineering, vocabulary updates or template refinement.

    - **No Autonomous Learning:** Heidi does not learn from user inputs or patient data, ensuring that bias or errors do not propagate or compound over time. All updates are manually engineered based on user feedback and based on QA/QC reports from our internal team of clinicians and engineers.

5. Can your AI scribe integrate seamlessly with electronic health record (EHR) systems commonly used in behavioral health, and what customization options are available to align with different clinical workflows, i.e. can clinicians review transcripts or create custom templates?

- Integrations are dependent on the specific EHR in question, however Heidi offers the flexibility for nearly any modern EHR to conveniently integrate Heidi into their platform via a variety of integration options (API and SDK Widget). Heidi also offers a Chrome Extension that is compatible with many Web-based EHRs for instances where the EHR is not integrated with Heidi.

- You can see it in action here -> **Chrome Extension Video** - Watch Video

- Heidi is known for its level of customizability and adeptness across any clinical workflow and multidisciplinary settings. This is primarily achieved through our powerful custom templates feature, in addition to complimentary features such as our 'context' tab discussed below. Clinicians can set up templates in a manner that not only determines structure of the output, but also the style and behavior based on natural language instructions that can be provided within a template in Heidi. We also dedicate template support for larger teams & organizations, and have a public template community where users from many specialties share their templates for many different use-cases.

- Additionally, users can leverage the 'context' tab available in sessions to provide Heidi with any additional context relevant to the clinical encounter (type, paste or upload), from real-time observations to historical notes. Heidi can leverage this contextual information to then create more comprehensive documentation, and alleviate the need for the clinician to manually incorporate that context into their clinical documentation.

- Lastly, transcripts are fully visible as part of the session, and users can review the session transcript at any time.

6. How long are transcripts and audio files stored?

- Transcripts generated during **live consultations** (no audio is ever stored by default) are stored **only as long as needed for the clinician to review, edit, and export them** to the medical record. Once exported, transcripts can be deleted by the user or set to expire automatically based on the clinic's data retention settings.

7. What are the pricing options for your product? Is your product designed for large health care systems or smaller practice settings?

- Heidi is designed for users in any clinical setting.

- We offer a truly **Free** tier allowing users to generate limitless notes using preset structures/ templates.

- In addition to our powerful Free Tier, we are one of the most affordable options in the market for our paid tiers, with prices ranging from as low as $66/month to $100/month (some options billed annually). Our paid tiers available suit clinical teams small and large, who are after a more extensive feature set (such as custom templates) ~ pricing depends on the feature-set and level of support chosen by users.

- We also work with some of the nation's largest health care systems across many specialties in both inpatient and outpatient care settings. This versatility is achieved through our custom

template capabilities and additional unique features native to Heidi, along with EHR integrations.

8. What if any studies have been done showing evidence for the efficacy and effectiveness of the tool?

- While we don't have any explicit studies testing Heidi specifically, there's some resources included below for to read through of use-cases of Heidi in a variety of mental health care settings:
    - **Psychiatry + Primary Care Clinic:** FQHC Hawse Health: 2 hours saved & 2 more patients seen per day with Heidi
    - UPenn trained Psychiatrist Dr. Tony Fernando uncovers the lift Heidi has had on his clinic
    - Why one of Indiana's largest multi-specialty behavioral health groups chose Heidi - Indiana Health Group