Nabla ambient AI scribe and Whisper hallucination concerns

IMPLEMENTED 2024 ONWARDS · HEALTH AND LIFE SCIENCES · FRANCE · BAND D · CLAIMED ONLY

VERSION 1.0 · PUBLISHED 17 NOVEMBER 2026 · UPDATED 17 NOVEMBER 2026 · FACTS CHECKED 17 NOVEMBER 2026

Nabla's ambient AI assistant drafts clinical notes from consultations for tens of thousands of clinicians. In 2024 the AP reported that the Whisper model it builds on can invent text and that the tool erased the audio needed to check transcripts; Nabla said clinicians review every note.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

AI-native company: built around AI from the start and attempting transformational change. It claims band A; MarLuc grades it on what is evidenced.

AI used: Ambient speech-to-text based on OpenAI's Whisper, retrained by Nabla, with a language model drafting clinical notes.

Implemented: 2024 onwards. Grade: band D, Task support, Claimed only.

Headline result, as reported: Used by more than 100,000 clinicians in over 190 healthcare organisations, by Nabla's account in July 2026; the AP reported hallucination risks in 2024

1 · The case

The facts, as reported by the 5 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Nabla. Clinical AI company with offices in France and the United States, per the AP, selling an ambient AI assistant, dictation and coding tools to health systems. Co-founded by Alex Lebrun, Delphine Groll and Martin Raison; Brian Manning became chief executive in July 2026. Canvas block: 6. Data and technology. Also 8. Governance and risk, 4. Operations and process.

Business challenge

The challenge. Clinicians spend time on note-taking and report writing; medical centres use speech-to-text tools to free them from it, the AP reports. Nabla says its mission is to restore the human connection at the heart of healthcare.

How it was addressed. Nabla's tool listens to the consultation, transcribes it with a model based on OpenAI's Whisper and fine-tuned on medical language, and drafts a structured note. Nabla says it spent 3 years and $5 million building a dataset of 7,000 hours of annotated medical audio to train its own model; that the transcript is not placed directly in the record; and that each note is split into atomic facts checked by a language model query against the transcript and patient context. Audio is not stored by default; clinicians can opt in per encounter where policy allows and the patient consents. In June 2025 Nabla said it would extend the assistant to coding and actions in the electronic health record. [1] [2] [3] [4] [5]

Who implemented it

OpenAI (Model or cloud provider): Whisper speech recognition model used as the base for Nabla's speech-to-text.
In-house (In-house): Nabla trained its own speech model and built the note-drafting product.

Technology (SEAM pillar E)

SEAM types: T5 Custom / fine-tuned models. Deployment: Not reported. Vendors and products: OpenAI, Whisper, Nabla ambient AI assistant.

Where it was built in: Listens to consultations and drafts notes that clinicians review and export to the record; Nabla integrates with Epic, Cerner, athenahealth, NextGen, Greenway and other electronic health records.

What changed

Before: clinicians wrote their own notes. After: the assistant drafts the note and the clinician who saw the patient reviews and edits it before export, by Nabla's account. Measured results: Nabla reports burnout symptoms down an average of 38% and documentation time cut by more than half; no independent figures. [1] [2] [3] [4] [5]

What it cost

Not reported

What it saved

Not reported in money. Nabla says it cuts documentation time by more than half and that burnout symptoms dropped by an average of 38%; these are Nabla's figures. [1] [2] [3] [4] [5]

Governance and control

Nabla says clinicians review and edit notes before export, a feedback tool flags discrepancies, and each model version is evaluated by professional medical scribes. The AP reported that erasing the audio made it impossible to compare transcripts with recordings, and that Nabla said its review requirement could change. No regulator or court action is reported.

S2 · Accountable owner · Not reported
Not reported.

A1 · Safe use · Claimed
Nabla says it does not store full recordings by default, for ethics and privacy reasons, and the product reminds users to review and edit notes before export. [2]

A2 · Human in the loop · Claimed
Nabla told the AP its tool requires providers to edit and approve notes, though that could change; its blog says the clinician who saw the patient reviews the note. [1] [2]

A3 · Monitoring and assurance · Claimed
Nabla says it monitors manual edits and user feedback and has each model version evaluated by medical scribes; the AP reported transcripts could not be compared with recordings because audio was erased. [1] [2]

A4 · Vendor AI · Claimed
The tool builds on OpenAI's Whisper, which OpenAI warns should not be used in high-risk domains, the AP reported; Nabla says it retrained the model on 7,000 hours of medical audio to suppress hallucinations. [1] [2]

A5 · Agent controls · Not reported
In 2025 Nabla said it would build agents that initiate actions in the electronic health record; controls on them are not reported. [4]

A6 · Standards and regulation · Not reported
No regulator or court action reported. Nabla says it is a member of the Coalition for Health AI. [4]

What worked, what did not, what is unproven

WORKED

– Rapid adoption: from over 30,000 clinicians in October 2024 to more than 100,000 in July 2026, by Nabla's account.

– Reduced documentation time and burnout, by Nabla's account.

DID NOT

– Researchers told the AP that Whisper, the base model, invents text, including a non-existent medication in one study example; no error in a Nabla note is reported.

– Transcripts could not be checked against the original recording because the tool erased the audio, the AP reported; a former OpenAI engineer said this removes the ground truth.

UNPROVEN

– Hallucination and error rates in Nabla's production system.

– Whether clinician review catches errors when notes are approved quickly.

– Effects on patient outcomes and record accuracy.

Timeline

Feb 2024 Careless Whisper study finds roughly 1% of Whisper transcriptions contain hallucinated phrases [3]

Oct 2024 AP reports Whisper hallucinations and that Nabla's tool, used by over 30,000 clinicians, erases the original audio [1]

Oct 2024 Nabla publishes how it uses Whisper: retrained model, fact checks, clinician review, opt-in audio storage [2]

Jun 2025 Nabla raises $70 million and plans agentic features that act in the electronic health record [4]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: Not applicable. Grade: D, Task support. Proof: Claimed only. Level of change: Task support.

Task support (Level 1, band D): clinicians do the same work and review AI-drafted notes. No independent measure of time saved or accuracy.

Task support: clinicians do the same work and review AI-drafted notes. Nabla is an AI-native company; its benefit figures are its own, and the independent sources concern risk in the base model.

The SEAM view and lessons learnt

The sources credit the uptake to demand for less documentation work; the benefit figures come from Nabla, so the AI's effect cannot be separated from other changes in client organisations. The AP's reporting shows the main safeguard is clinician review, which carries more weight when audio is not kept. Before copying, a leadership team should ask how errors in AI-drafted notes will be found and measured, and whether recordings are kept long enough to audit a sample.

1. Decide how AI-drafted clinical notes will be checked before rollout.
2. Keep a consented way to verify transcripts for sampled audits.
3. Ask suppliers for measured error rates in your own setting.

Evidence

Score 1 of 4. Mixed: usage and benefit figures are Nabla's own; the AP investigation and the Careless Whisper study are independent but concern the base model, not a measured error rate for Nabla's product.

Evaluation method. Self-reported company figures, set against independent press reporting and an academic study of the base model.

Sources disagree. In October 2024 the AP reported 40 health systems and an estimated 7 million visits; Nabla's blog that week said over 85 healthcare organisations, more than 45,000 clinicians and 9 million encounters in 24 months. The AP reported that the tool erases the original audio; Nabla says clinicians can opt in to storing audio per encounter, with patient consent.

Status. Ongoing: Nabla reported more than 190 healthcare organisations and 100,000 clinicians in July 2026.

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/nabla-ambient-ai-scribe-whisper

[1] Anchorage Daily News (Associated Press): Researchers say an AI-powered transcription tool used in hospitals invents things no one ever said (October 2024)
https://www.adn.com/nation-world/2024/10/26/researchers-say-an-ai-powered-transcription-tool-used-in-hospitals-invents-things-no-one-ever-said/

[2] Nabla: How Nabla uses Whisper (October 2024)
https://www.nabla.com/blog/how-nabla-uses-whisper

[3] arXiv: Koenecke and colleagues, Careless Whisper: Speech-to-Text Hallucination Harms
https://arxiv.org/abs/2402.08021

[4] Nabla: Nabla raises $70M Series C (June 2025)
https://nabla.com/press-release/nabla-raises-70m-series-c-to-deliver-agentic-ai-to-the-heart-of-clinical-workflows-bringing-total-funding-to-120m

[5] Nabla: Nabla appoints Brian Manning as Chief Executive Officer (July 2026)
https://nabla.com/press-release/nabla-appoints-brian-manning-as-chief-executive-officer-to-lead-next-phase-of-growth

Nabla website: https://www.nabla.com

Version 1.0 · Published 17 November 2026 · Updated 17 November 2026 · Facts checked against sources 17 November 2026 · Downloaded 26 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/nabla-ambient-ai-scribe-whisper

VERSION HISTORY

v1.0 17 November 2026 First published