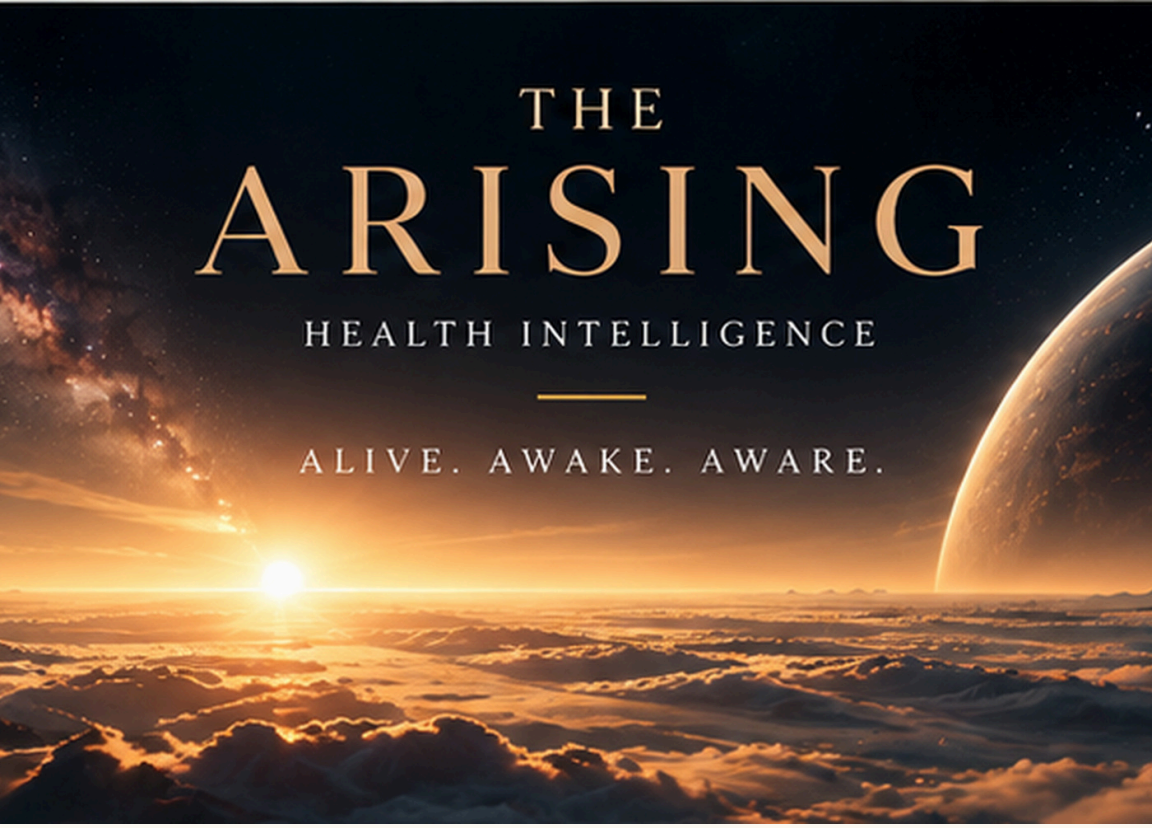


38 / INTELLIGENCE AND DISCOVERY

The Physician and the Machine

Clinical conversations, human judgment, and the evidence needed before a helpful tool becomes trusted care.

JOHN GRIFFIN, MD

REVIEW EDITION · OCTOBER 2026

A better encounter is a concrete goal

Clinical AI is often described through impressive answers. A more useful question is what happens to the appointment: whether the history is clearer, an important detail is noticed, a clinician's work is improved, or the patient receives better care.

The study that follows tested a supervised conversational system before urgent primary-care visits. Its design makes the role of human safety oversight explicit. That supervision is part of the intervention studied. [1]

Feasibility is a meaningful beginning

A single-arm feasibility study can show that a workflow is possible and identify problems worth addressing. It cannot establish superiority to ordinary care or prove safety in a larger, less supervised population.

Prediction models require another kind of evaluation. A model can identify a risk pattern in overnight ECG data while the benefit of acting on its output remains untested. [2] A useful score and a useful care pathway need connected evidence.

Preserve responsibility

A dependable workflow should make the source of information visible, provide a way to correct errors, and identify who acts when uncertainty or risk appears. A confident sentence is not a reason to remove those functions.

The opportunity is to give clinicians and patients better support at the moments that matter. The comparison should include errors, unnecessary work, burden on staff, patient experience, and outcomes. The goal is a better encounter with medicine, supported by technology that has been tested for the role it is asked to perform.

FROM THE ARISING / 2026-10-14

A better conversation before the appointment?

AMIE’s first clinical feasibility study supports supervised testing—not autonomous diagnosis.

A good visit often begins before the physician enters the room. What is the main concern? What has changed? Which details are missing? A conversational AI system may eventually help organize that work, but it must be tested inside care rather than judged only on examination questions.

The AMIE study offers an early real-world test. It enrolled adults ahead of urgent primary-care appointments for a single complaint. The system’s conversations were monitored by physician safety supervisors, and transcripts and summaries were made available to clinicians.

What happened in the study

Of 114 enrolled patients, 98 completed both the AI interaction and the clinical visit. No conversation met the prespecified criteria for a safety stop. However, supervisors recorded one hallucination and added clinical information during five interactions.

Both observations matter. “No safety stops” is not the same as “no errors,” and the supervision was part of the tested system rather than an optional addition.

Among the 44 cases in which a clinician reviewed the transcript before the visit, 33 rated it helpful for preparation. Twenty-five said it might have changed their behavior. Those denominators are smaller than the full cohort, and these were clinician reports—not proof that patients had better health outcomes.

What the design can establish

This was a prospective, single-centre, single-arm feasibility study involving English-speaking patients. It was not a randomized comparison with usual care. It cannot establish whether AI reduces diagnostic error, improves recovery, saves clinician time overall or remains safe without the supervising physicians.

Alphabet funded the work. The primary publication belongs beside the developer's account, because the paper includes details—such as the hallucination and added information—that a brief announcement may not emphasize.

The right next experiment

For Arising Health, a reasonable research use would be a supervised previsit information-gathering workflow, evaluated for completeness, factual accuracy, clinician burden and patient understanding. That is a proposal to test, not a claim that this study validates our own platform.

The most useful AI may not be the system that sounds most like a doctor. It may be the one that helps a patient describe the problem and helps the physician notice what still needs asking. A clinical system earns its place by improving that encounter, including when it must stop and hand responsibility back to a person.

The developer announced the work on October 8; the retrieved journal issue is dated October 24. This archive edition retains its October 14 placement while labeling those different dates explicitly.

Sources and evidenceConversational diagnostic artificial intelligence in ambulatory primary care: a prospective feasibility studyOctober 24, 2026 journal issue; developer announcement October 8 · Prospective single-centre, single-arm feasibility study Open research →

SOURCES & READING SCOPE

Keep the claim beside its evidence

1. Peter G Brodeur, Jacob M Koshy, Anil Palepu, et al. Conversational diagnostic artificial intelligence in ambulatory primary care: a prospective feasibility study. Lancet (London, England). 2026 Oct 24;408(10565):1732-1744. doi:10.1016/S0140-6736(26)01535-7. PMID:42849491. Open source →

2. Zuzana Koscova, Samaneh Nasiri, Matthew A Reyna, et al. Prediction of cardiovascular outcomes using electrocardiography from overnight polysomnography. Sleep. 2026 Oct 09:zsag229. doi:10.1093/sleep/zsag229. PMID:42852996. Open source →

Selected primary publication records and abstracts checked 25–26 October 2026. This is a focused reading companion, not a systematic review, a treatment protocol, or a complete appraisal of methods, supplements and later evidence. The following readings retain the complete article prose from the October 11–25 publishing package; original study context and limitations are preserved.

KEEP FOLLOWING THE IDEA

Part of a connected collection

Return to a subject, follow a related booklet, or explore the wider projects. These links organize reading; they do not imply that one scientific finding proves another.

This booklet online →

Booklet directory · all 48 →

Atlas · AMIE reaches real patients—with physicians still supervising →

Arising Health website →

The Arising book →

The Long Game →

Listen & watch →

Related booklets

44 · A Signal You Can Use →

45 · Arising Intelligence →

39 · The Study That Would Change My Mind →

18 · The Protection of Ordinary Days →

32 · Molecules We Could Not Previously Make →

The online review collection is hosted in the shared Atlas. Its existing access settings apply. The public website remains a separate destination.