

PSA Morning Briefing

Friday, September 11, 2026

Good morning, JD. Today's strongest message is that underserved people are already turning to AI when healthcare is unavailable, expensive, or intimidating. PSA's opportunity is to make that use safer while investing in the human support that technology cannot replace.

Only three sufficiently relevant developments surfaced from the last 48 hours. I've supplemented them with two clearly dated, practical resources rather than manufacture novelty.


1. Health information is rapidly moving into AI—but trust remains thin



Figure 1

One in Three Adults Have Used AI for Health Information and Advice in the Past Year


Percent who say they have sought information or advice about their...from each of the following sources in the past year:



Note: See topline for full question wording.


Source: KFF Tracking Poll on Health Information and Trust (Feb. 24-Mar. 2, 2026)

KFF



Many Adults Who Used AI for Health Information Did Not Later Follow Up With a Doctor

After using an AI tool for information or advice related to your...did you follow up with a doctor or other health professional?



Note: Among adults who have used AI for information or advice related to their physical health or mental health, respectively, in the past year. See topline for full question wording.

Source: KFF Tracking Poll on Health Information and Trust (Feb. 24-Mar. 2, 2026)

KFF



SymptomAI: Towards a Conversational AI Agent for Everyday Symptom Assessment

Academic poster summary of Breda et al. | Google Research & Google DeepMind | 2026



KEY TAKEAWAY

Conversational AI that actively interviews users about symptoms can outperform clinician-generated differential diagnoses from the same transcript and can enable large-scale wearable biomarker discovery.

1 BACKGROUND & MOTIVATION



Large language models perform well on curated medical vignettes, but real-world everyday symptom assessment by laypeople is much harder.



SymptomAI was designed to test whether a conversational AI agent can conduct symptom interviews, provide differential diagnoses (DDx), and support population-scale health analysis.

2 STUDY DESIGN



Participant in Fitbit App

9-month real-world deployment through the Fitbit app



Conversational AI Interview

N = 13,917 participants



Clinical Review & Evaluation

N = 517 cases underwent expert clinical evaluation



Self-Reported Diagnosis (Optional)

N = 1,228 participants self-reported a clinician-provided diagnosis



Wearable Data Collection

>500,000 wearable-data days analyzed



Population-Scale Health Analysis

Nearly 400 unique conditions examined



Auxiliary General-Population Study

N = 1,509 participants

3 PROMPTING STRATEGIES



1 Base User-guided conversation



2 Fixed Canonical Fixed history-taking questions



3 Flexible Canonical Canonical questions with flexible order



4 Dynamic Live Fully agent follow-up with live DDx updates




5 Dynamic Final Fully agent follow-up with final DDx only



Strategies that actively elicited more information significantly outperformed the user-guided base condition.

4 MAIN RESULTS

Top-5 Differential Diagnosis (DDx) Accuracy



Information-eliciting strategies achieved an average 27.34% higher accuracy than the base user-guided strategy.



Dynamic agentic strategies performed comparably to canonical history-taking strategies: 71.4% vs 75.6%; $p = 0.155$.



Blinded clinicians ranked SymptomAI as the best DDx in over 50% of cases.



Auxiliary study showed similar performance in a broader US population: 75.2% top-5 accuracy vs 80.0% in the Fitbit study.

Robustness: When Information or Confidence is Low



SymptomAI particularly outperformed clinicians when conversations contained limited information or when clinicians were less confident.




Accuracy by Prompting Strategy (Top-5)



All information-eliciting strategies outperformed the user-guided base condition ($p < 0.001$).

Average improvement over base: +27.34%

Comparison of Key Strategy Families



5 WEARABLE BIOMARKER INSIGHTS



SymptomAI labels enabled phenome-wide analyses across nearly 400 conditions.



Strong associations were found between acute infections and physiological shifts.




Influenza showed strong biomarker associations with odds ratios greater than 7.



In respiratory infection cohorts, resting heart rate and respiration tended to increase near symptom onset, while HRV and daily activity tended to decrease.

Biosignal Trends Around Symptom-Report Day (0)



6 IMPLICATIONS



- Declared symptom interviews improve AI diagnostic performance.
- Conversational AI may expand access to symptom assessment.
- AI-generated phenotyping can unlock large-scale digital biomarker research.
- Wearables may help trigger proactive symptom-check conversations.

7 LIMITATIONS



- Many diagnoses were self-reported.
- Clinical evaluation used retrospective transcript review.
- Research system only; not a substitute for medical care.
- Fitbit-based recruitment may not perfectly represent every population.

8 CONCLUSION



SymptomAI demonstrates that agent-led, multi-turn symptom interviewing can improve everyday differential diagnosis performance and create scalable links between conversational health data and wearable physiology.

Event and organization imagery; not a photograph of survey participants.

The Consumer Technology Association announced research on September 9 finding that **four in ten consumers turn to AI or social media for health information**. The publicly accessible [CTA release listing](#) provides the headline finding but not enough methodology to judge sampling, question wording, or subgroup results. CTA is also a technology-industry association, so this should be treated as market research—not independent clinical evidence.

Still, the direction agrees with earlier independent polling. A West Health–Gallup poll found approximately one-quarter of American adults had recently used AI for health information. Cost, after-hours access, embarrassment, and previous experiences of being dismissed were among the reasons people used it. Only about one-third trusted its accuracy, while most continued consulting healthcare professionals. [Associated Press summary of the polling](#).

Why PSA should care: People in southern New Mexico will not wait for a formally approved application before asking AI about symptoms, medications, or test results. Simply advising them never to do so is unlikely to work.

42's take: PSA should teach **safe use rather than silent use**: remove identifying information, ask for sources, verify medication advice, and recognize questions requiring a clinician or poison-control center. AI can help someone prepare for care; it should not impersonate care.

Practical follow-up: Pilot a bilingual, 45-minute “Ask AI About Health—Safely” workshop using fictional cases. Measure whether participants can identify hallucinations, privacy risks, and urgent symptoms afterward.

Near-term prospect: Immediately feasible at low cost, provided the workshop uses no real patient information and a qualified health professional reviews the curriculum.

2. Real language-access pilots point toward AI-assisted documents—not autonomous interpreting

Official report-related imagery from the California Health Care Foundation.

A March 20 California Health Care Foundation report provides one of the most useful current maps of AI language-access pilots. Healthcare organizations are concentrating on comparatively lower-risk work: pre-visit instructions, reminders, patient education, and written post-discharge information. They are generally using AI to support qualified staff, not replace interpreters. [Issue brief and report](#).

One case describes Children's Hospital Los Angeles using an AI-assisted workflow for written Spanish discharge instructions. According to staff interviews, preparation time reportedly fell from roughly two hours to 5–10 minutes, with multilingual professionals reviewing the output. That is an encouraging operational result, but it is **not a controlled study of comprehension, medication safety, or readmission rates**.

The report estimates conventional telephone interpretation at roughly \$2–\$5 per minute, on-site interpretation at \$45–\$150 or more per hour, and document translation at \$0.10–\$0.30 per word. It does not establish that an AI vendor will have a lower total cost once licensing, integration, quality assurance, and staff review are included.

Why PSA should care: Spanish-language after-visit materials are a realistic southern New Mexico use case. Consent, diagnosis, medication changes, and emotionally sensitive conversations should continue to involve qualified human interpreters.

42's take: The best first pilot is a **document-production tool with a named human approver**, not live autonomous medical interpreting.

Practical follow-up: Ask a clinical partner to select five frequently used, non-urgent English handouts. Compare professional translation alone with AI-assisted translation plus professional review for time, cost, reading level, errors, and patient comprehension.

Near-term prospect: Technically feasible now. Clinical partnership, vendor privacy terms, and qualified review are non-negotiable.

3. Multilingual education research offers a useful teacher-first AI design

KHUNG THÍ ĐIỂM GIÁO DỤC TRÍ TUỆ NHÂN TẠO (AI) CHO HỌC SINH PHỔ THÔNG

Quyết định số 3439/QĐ-BGDĐT ngày 15-12-2025 của Bộ Giáo dục và Đào tạo

TƯ DUY LẤY CON NGƯỜI LÀM TRUNG TÂM

Mục đích AI:
Phục vụ con người.
Kỹ năng:
Xác định nhu cầu và đánh giá
giải pháp AI.

THIẾT KẾ HỆ THỐNG AI

Năng lực từ sử dụng đến
tự tạo hệ thống đơn giản.
Tập trung kỹ năng
giải quyết vấn đề.



ĐẠO ĐỨC AI

Nhận diện & phân tích
vấn đề đạo đức, pháp lý.
Hướng tới sử dụng AI
trách nhiệm.

CÁC KỸ THUẬT VÀ ỨNG DỤNG AI

Kiến thức cơ bản về
cách AI hoạt động.
Kỹ năng sử dụng
các công cụ AI trong thực tế.

Overview



Out of three students grade 1-12 in Vietnam, there is one student uses VioEdu



18.5M

Total users
(2019-2023)

6M

Mobile app downloads
(2019- 2023)

1.8M

Active users/month
(2023)

#8


Keyword of the year on Google
Search (2021, Vietnam)

Top Rank

Educational app
(Appstore, Vietnam)

40+ /63

Provinces corporate



Giáo viên thời đại AI mới > Khám phá GenAI - Hiểu và ứng d... > 1.4. Đánh giá tổng kết

1.4. Đánh giá tổng kết

Câu 1

Câu 2

Câu 1

1 điểm

Trong 4 lựa chọn sau, #Ngữ cảnh phù hợp nhất cho tiết Ngữ văn 8?

☒ A

 "Bạn là nhà văn, hãy viết truyện ngắn..."

☐ B

 "Bạn là GV Ngữ văn 8, dạy văn bản Chiếc lá cuối cùng, lớp có 2 HS ELL, mục tiêu: nhận biết - thông hiểu - vận dụng."

☐ C

 "Bạn là học sinh lớp 8"

☐ D

 "Bạn là chuyên gia ngôn ngữ học"

Câu trả lời chưa chính xác. Bạn có thể nhận gợi ý bằng cách hỏi OctoAI.

Thử lại

Hỏi Bạch Tuộc Biết Tuốt

Số lần trả lời: 1/3

Câu 2

1 điểm

Trong các câu lệnh sau, câu nào chưa có đủ phần "dữ liệu đầu vào"?

☐ A

 "Hãy sửa lỗi chính tả trong đoạn văn sau: 'Hôm nay em đi học muộn, cô giáo phạt em trực nhật lớp.'"

☐ B

 "Dưới đây là bài làm học sinh. Em hãy gợi ý 3 cách cải thiện."

☐ C

 "Viết 5 câu hỏi đọc hiểu cho văn bản lớp 6."

Bạch Tuộc Biết Tuốt

Câu trả lời này chưa đúng rồi. Bạn có thể cho mình biết vì sao bạn chọn đáp án đó không?

Tôi nghĩ mình chỉ cần cho AI biết đang đóng vai gì và nên làm gì là đủ rồi

Bạn ơi, "dữ liệu đầu vào" chính là thông tin/yêu cầu mà bạn cung cấp cho AI để AI hiểu và tạo ra kết quả đúng ý bạn. Nếu đầu vào mơ hồ hoặc thiếu thông tin, AI sẽ cho kết quả không phù hợp

Tôi hiểu rồi. Cảm ơn bạn!

Bạch Tuộc Biết Tuốt có thể cung cấp thông tin không chính xác. Thấy/Có hãy xác thực câu trả lời. Điều khoản | Chính sách

Nhập tin nhắn...

zoom Workplace

PowerPoint Slide Show - [250927-TậpHuấn AHK1] Bạn đang xem màn hình của: Cẩn Thơ_THCS_AI_1 (1985948) Xem tùy chọn

viettel

2.2.2. THCS - AI Technology

Mạch kiến thức

5. Huấn luyện mô hình (quy trình thu thập - huấn luyện - thử nghiệm) L6-L7

6. Lập trình điều khiển mô hình trong ứng dụng trực quan L6-L7

7. Học máy (Machine Learning pipeline) L7

8. Học có giám sát (Supervised learning: phân loại, dự đoán) L7

9. Rule-based (chatbot theo kịch bản) L8

10. Hội thoại, lập trình Chatbot L8

11. Luồng và trạng thái trong hội thoại (state-flow) L8-L9

Vị trí trong hệ tri thức AI

Làm quen Machine Learning cơ bản (supervised learning).

Thấy rõ sự khác biệt giữa AI rule-based (ký hiệu) và AI học máy (dữ liệu).

Hiểu về tương tác người - máy (NLP ứng dụng, chatbot).

Chuẩn mực

Phê duyệt

Hiện đại

Slide 16 of 56

Type here to search

8:26 AM 9/21/2025

lil Cẩn Thơ_THCS_AI_1 (1985948)

Tập huấn giáo viên môn Trí tuệ nhân tạo_27/09/2025



Research-system or paper-related imagery; no suitable documentary image of the study participants was confirmed.

A June 2026 preprint describes ConnectED, a Vietnamese-language AI system that links curriculum-based lesson planning, student practice, and teacher review. Instead of giving a chatbot an unrestricted teaching role, its workflow places teacher-validation checkpoints throughout lesson creation. [Original preprint](#).

The underlying model answered **87.02% of 3,119 questions** drawn from Vietnam's 2025 national high-school examination, reportedly outperforming its comparison model by 6.1 percentage points. Surveys included 18 teachers and 214 students. Lesson-preparation time was reported to decline from three or four hours to approximately 30–45 minutes.

These are promising early results, but the limitations are important. The paper is a preprint, the teacher sample is small, and exam-question accuracy is not proof of lasting student

learning. The study does not establish effectiveness for Spanish-English learners, American curricula, students with disabilities, or low-bandwidth households. ConnectED also appears to be a research system, not a confirmed off-the-shelf product PSA can purchase.


Why PSA should care: The transferable insight is the workflow: anchor AI to an approved curriculum, constrain its tasks, require educator review, and use student performance to revise lessons. That is much safer than asking a general chatbot to “teach algebra.”


42’s take: PSA could reproduce the design principles without adopting this particular system. A bilingual educator might use AI to generate alternate explanations and practice questions, approve them, and then measure learning away from the tool.

Practical follow-up: Run a four-week Spanish-English numeracy pilot using one approved lesson sequence. Record preparation time, correction rates, independent post-tests, and the number of culturally or linguistically inappropriate outputs.

Near-term prospect: Pilotable with existing tools; local evidence would still be needed before expansion.

4. A national higher-education AI institute begins its new cycle





Official program and participating-campus imagery.

The Association of American Colleges and Universities began its 2026–27 **Institute on AI, Pedagogy, and the Curriculum** on September 10. The online, team-oriented program runs through April 6, 2027. AAC&U says approximately 330 institutional teams have participated across its institute cycles. [Official institute listing](#) and [AAC&U program webinar](#).

The program's importance is organizational rather than technological. Colleges commonly have isolated instructors experimenting with AI while admissions, disability services, libraries, faculty, and students operate under different assumptions. A multi-month

institute gives campus teams time to coordinate curriculum, assessment, policy, and professional development.

This launch is evidence that institutions are investing in structured planning; it is not evidence that participation improves graduation, employment, or learning outcomes. I also did not confirm an open place for PSA or a published participation price from accessible program information.

Why PSA should care: Doña Ana Community College, New Mexico State University, adult-education providers, and community organizations face the same coordination problem. Low-income and bilingual students are most vulnerable when rules are unclear or when instructors assume everyone owns a capable device and reliable connection.

42's take: PSA does not need a seven-month national program to borrow its strongest feature: put educators, learners, community representatives, accessibility specialists, and administrators at the same table. Students should help define acceptable AI use rather than merely receive a finished policy.

Practical follow-up: Convene a two-session “southern New Mexico mini-institute.” Session one identifies three learner problems; session two evaluates one proposed pilot against learning, language access, accessibility, privacy, and complete cost.

Near-term prospect: Convening is immediately realistic. Institutional policy changes would likely take an academic term or longer.