ALLIANCE FOR HEALTH POLICY

THE ALLIANCE SIGNATURE SERIES

Health Data Privacy in the AI Era: Navigating Federal and State Protections and Emerging Risks

Public Congressional Briefing
Thursday, October 22, 2026

Moderator & Panelists



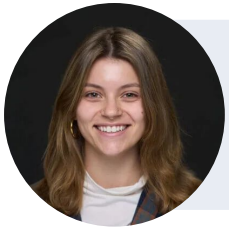
Jodi Daniel, J.D., MPH
MODERATOR
Partner
Wilson Sonsini Goodrich & Rosati



Deven McGraw, J.D., MPH
Chief Regulatory and Privacy Officer
Citizen Health



Randi Seigel, J.D.
Partner
Manatt, Phelps & Phillips, LLP



Emma Beavins
Marketing and Content Lead
Vega Health



Kev Coleman
Director, Health Care AI Initiative
Paragon Health Institute



Grace Cordovano, Ph.D., BCPA
Founder, Enlightening Results
Co-Founder, Unblock Health



Kev Coleman

Director, Health Care AI Initiative
Paragon Health Institute

Health Data Privacy In the AI Era:

TECHNICAL CONSIDERATIONS

Kev Coleman

PARAGON HEALTH INSTITUTE

Agenda

01 AI DIVERSITY

02 GEN AI

03 AGENTIC AI

04 OFF LABEL AI

05 WHAT'S AHEAD

AI Diversity



- AI is a collection of different programming architectures, **not a singular type of software**
- Different types of AI pose different privacy risks to health data
- AI data processing may be performed locally at a health facility, in the cloud, or a hybrid of the two
- AI in patient-facing health systems spans three major contexts:
  - Front office: Appointment scheduling, after-hours chatbots, verification of insurance & copays, etc.
  - Clinical: Ambient scribing, medical image interpretation, CDS, prior authorization requests, readmission prediction, etc.
  - Back office: Automated medical service coding, claims processing, claim denial response, ERP, etc.
  - Patient data may travel back and forth across all three domains



AI Diversity

As the later off-label AI slide will highlight, health data privacy in the AI era is more than a matter of technical cybersecurity. It is deeply entangled in issues such as:

- Internal controls (i.e., privacy policies and procedures working alongside technical defenses)
- Staff behavior
- Patient behavior

Gen AI

- Generative AI, or Gen AI, is a type of AI that creates original outputs such as text, images, video, audio, and computer programming code
- Large language models (LLMs) are a popular form of gen AI in healthcare, but generative adversarial networks (GANs) are also used less commonly
- Popular gen AI systems include ChatGPT, Claude, Grok, and Gemini
- Gen AI presents several privacy risks:
  - Memorization: user inputs or training data can, in some cases, be memorized
    - Stanford/Yale book test
  - Future training: Gen AI may update training based on user inputs unless those inputs have been excluded by an opt-out choice
  - Re-identification: Gen AI may, under various conditions, be able to reconstruct the identity associated with de-identified personal data
  - Storage: Like other software systems, data retained (even temporarily) in an AI system may be breached

Agentic AI

- Agentic AI refers to AI (usually built on an LLM foundation) that reasons how to accomplish a user-provided goal & executes the resulting multi-step workflow without little to no human supervision
- Agentic workflows may connect to multiple hospital computer systems (e.g. EHRs), extract information (including PHI), interpret, and then produce reports, generate invoices, order supplies, create surgery schedules, etc.
- With respect to data privacy, the most important aspects of agentic AI concern:
  - Cross-system access
  - Data integration and synthesis
  - Non-approved agents created by health system staff
  - Rogue agents & hacking (e.g., Australia Medicare Stats portal)

Off-Label AI

- Off-label AI, otherwise known as “Ghost AI,” refers to the unofficial use of AI by medical staff, including the use of tools that have not been approved for medical use
- Patient data may be entered into a commercial AI system without the knowledge of the medical facility
- Unofficial use of AI means that there is no governance over data inputs, AI applications used, security guarantees for data in motion and rest
- Off-label AI extends beyond healthcare professionals to patients
  - Patients may upload medical docs to a foundation model (e.g., ChatGPT) to obtain a more personalized medical answer
  - PHI may be used for model training unless the consumer explicitly opts out
  - PHI disclosures by a patient can complicate matters when assessing the source of a data breach

Off-Label AI

- In early 2026 Wolters Kluwer Health announced findings indicating more than 40% of medical workers & administrators surveyed know colleagues who are using ghost AI. Nearly 20% acknowledged using ghost AI themselves
- Off-label AI poses the risk of “data leakage.” Data leakage is when protected health information (PHI) and/or personally identifiable information moves from an authorized environment into an unauthorized environment that does not have the same privacy controls and regulatory protections

What's Ahead

- Agentic AI will continue to proliferate across healthcare
- Phishing will increasingly leverage LLMs to scrape public healthcare provider information to craft (at scale) highly tailored emails, texts, and possibly voicemails using voice cloning
- Continuous learning will emerge as a force that transforms healthcare system operations
- AI development will continue to exert pressure on healthcare systems for interoperability and the sharing of de-identified patient data

Kev Coleman

Director of AI Healthcare Initiative
Paragon Health Institute
Paragon institute.org/initiatives/health-care-ai/



Deven McGraw, J.D., MPH

Chief Regulatory and Privacy Officer
Citizen Health

HEALTH AI PRIVACY FOUNDATIONS · PART 1

Mapping the federal health privacy landscape

Deven McGraw· October 22, 2026

Health data now comes from far beyond the clinic

Inside traditional care	Outside traditional care
Usually within HIPAA	Often originates outside HIPAA
Medical records and clinical notes	Wearable and fitness tracker data
Claims and billing data	Health, fertility and mental health apps
Lab results and imaging	Search, purchase and location history
Prescription and pharmacy records	Conversations with AI chatbots and assistants

Everyday data can also reveal health conditions through inference.

HIPAA* protects information held by specific entities

Who is covered

Health plans

Health care clearinghouses

Providers who bill or transact electronically

Business associates handling PHI for them

What the rules require

Privacy Rule: limits uses and disclosures, grants access rights rights

Security Rule: safeguards for electronic PHI

Breach Notification Rule: notice to people, HHS and sometimes media

PHI is identifiable health information held by a covered entity or business associate. Data meeting HIPAA's de-identification standards falls outside of HIPAA's coverage. (*Relevant HITECH provisions now part of HIPAA)

Two ways to deidentify under HIPAA

Safe Harbor

45 CFR 164.514(b)(2)

Remove 18 categories of identifiers of the person

Identifiers required to be removed include names, most dates, detailed geography, contact details, record numbers and photos

No actual knowledge that what remains could identify someone

Expert Determination

45 CFR 164.514(b)(1)

A qualified expert applies statistical or scientific methods

Concludes the risk of identifying anyone is very small, and documents why

Can keep useful detail, such as dates; considers context context (recipient & what data they can access to re-re-identify

Other federal authorities (HIPAA, where it applies, is a floor)

Authority	What it covers
42 CFR Part 2	Substance use disorder records, now more aligned with HIPAA (Feb 16, 2026)
GINA	Genetic information in health insurance and employment decisions
FTC Act, Section 5	Unfair or deceptive data practices by most companies, including apps
FTC Health Breach Notification Rule	"Breaches" of health data held by non-HIPAA apps and record vendors (breach = use or disclosure without the individual's authorization)
21st Century Cures Act	Information blocking rules (creates expectations for sharing data when the law allows such sharing, including with apps patients choose)

Three categories of data, three different rule sets

Category	Main rules	What it means for AI
Protected health information or PHI	HIPAA, plus Part 2 for substance use disorder records	AI uses need a permitted purpose, minimum necessary limits and business associate agreements
Consumer health data	FTC Act, the Health Breach Notification Rule and state laws	Training and sharing depend largely on privacy promises and consent
Deidentified data	Outside HIPAA once Safe Harbor or Expert Determination is met	Can be pooled and used for training; linkage can raise reidentification risk



Randi Seigel, J.D.

Partner

Manatt, Phelps & Phillips, LLP



State and Federal Privacy Landscape

Presented by: Randi Seigel, Partner
Manatt, Phelps & Phillips, LLP

HIPAA Is Only the Starting Point

BROADER COVERAGE

State laws reach entities HIPAA does not.

Payers, website operators, digital health apps — and soon, AI chatbots.



TEXAS | Health & Safety Code Ch. 181

- Covers **anyone who handles PHI** — for profit, nonprofit, or pro bono — with actual or constructive knowledge
- Expressly names **business associates, payers, government, IT firms, schools, researchers, facilities, and website operators**
- Likely reaches anyone who **comes into possession of PHI**



CALIFORNIA | CMIA (Civ. Code § 56)

- **Digital health businesses are "providers"** (§ 56.06): health-info apps, software and devices, plus mental-health and reproductive/sexual-health digital services
- **AB 1979 adds AI health chatbots** — consumer-facing deployers covered from **Jan. 1, 2027** (signed Sept. 30, 2026)

Sources: California Medical Information Act (Cal. Civ. Code § 56 et seq.); Tex. Health & Safety Code ch. 181

States Add Stricter Rules for Sensitive Clinical Information

MENTAL AND BEHAVIORAL HEALTH

Protects the therapeutic relationship itself.



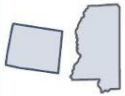
COLORADO | Mental-health professionals

No disclosure of client communications without consent, subject to enumerated exceptions



MASSACHUSETTS | Licensed psychologists

Communications and information from services are confidential, subject to specified exceptions



MISSISSIPPI & WYOMING | Mental-health professionals

No disclosure of professionally obtained information except under statutory exceptions

SUBSTANCE USE DISORDER

Protects SUD treatment records and programs.



NEW YORK | SUD programs

OASAS-certified or funded programs may not disclose patient records, subject to federal and state exceptions



LOUISIANA | Addiction professionals

Credentialed professionals may not disclose client information, subject to enumerated exceptions



MASSACHUSETTS | Treatment facilities

Covered treatment facilities may not disclose patient records, subject to specified exceptions

Other categories with heightened protection include reproductive health, HIV/STI status, and genetic information.

Sources: Colo. Rev. Stat. § 12-245-220; Miss. Code Ann. § 73-30-17; Wyo. Stat. Ann. § 33-38-109; 14 NYCRR Part 815; La. Rev. Stat. § 37:3390.4; Mass. Gen. Laws ch. 111E, § 18.

23 States Have Enacted Comprehensive Privacy Laws



19

Laws in effect

California was first (2020); Indiana, Kentucky and Rhode Island are newest (Jan. 1, 2026).

4

Enacted in 2026, not yet effective

Oklahoma and Louisiana (Jan. 1, 2027); Alabama (May 1, 2027); Vermont (Jan. 1, 2028).

- In effect
- Enacted, not yet effective
- Florida & Washington: outside IAPP's scope
- No comprehensive law

Sources: IAPP, US State Privacy Legislation Tracker (state law map, current as of Sept. 2026), iapp.org/resources/article/us-state-privacy-legislation-tracker. IAPP counts only comprehensive laws and excludes Florida's Digital Bill of Rights and Washington's My Health My Data Act.

Consumer Privacy Laws Reach Health Information Outside Traditional Health Care

WASHINGTON | MY HEALTH MY DATA

Sweeping consumer-health-data law reaching any business that collects health data from Washington consumers, including data that can infer a health condition, with no revenue threshold.

Requires separate consent to collect and to share, a signed authorization to sell, and a ban on geofencing within 2,000 feet of health care facilities.

MARYLAND | MODPA

Among the most restrictive state laws, with low thresholds (35,000 Maryland residents) and narrow exemptions.

Sensitive data, including health data, may be collected, processed or shared only when strictly necessary to provide a requested product or service, and its sale is banned outright. Data of consumers the controller knew or should have known are under 18 cannot be sold or used for targeted ads. A data-level HIPAA exemption leaves non-PHI in scope.

CALIFORNIA | CCPA / CPRA

Broad consumer privacy law treating health and neural data as sensitive personal information.

Consumer rights address collection, access, deletion, correction, sale or sharing and certain sensitive-data uses. Regulations effective Jan. 1, 2026 add risk assessments and automated decision-making technology rules.

Sources: Wash. Rev. Code ch. 19.373 (eff. Mar. 31, 2024); Md. Code Ann., Com. Law §§ 14-4701–14-4714 (eff. Oct. 1, 2025); Cal. Civ. Code § 1798.100 et seq. (CCPA eff. Jan. 1, 2020; CPRA amendments eff. Jan. 1, 2023; as amended by SB 923, signed Sept. 27, 2026, eff. Jan. 1, 2027; AB 1542 sensitive-data sale ban vetoed Sept. 27, 2026); Manatt: Washington's My Health My Data Act (Apr. 17, 2023); Maryland MODPA Now in Effect (Oct. 2, 2025); Cal. Att'y Gen.: CCPA (oag.ca.gov/privacy/ccpa).

Minors Receive Greater Protection Under Some Consumer Privacy Laws

MINORS

Age, knowledge and consent rules vary by state.

States differ on the age threshold, knowledge standard, who may consent, targeted ads, sale or sharing, profiling, retention, geolocation and engagement design.



CALIFORNIA | Under 16

A business that knows (or willfully disregards) that a consumer is under 16 may not sell or share their data without **opt-in authorization**: from the consumer if 13–15, or a parent or guardian if under 13.



MONTANA | Under 18

Duty of reasonable care to avoid heightened risk of harm to known minors; consent required for targeted ads, sale or profiling of their data.



COLORADO | Known minors

Adds **reasonable-care and assessment duties** and requires consent for targeted ads, sale, certain profiling, secondary use, extended retention and precise geolocation.

Sources: Cal. Civ. Code § 1798.120(c) (CCPA); Mont. Code Ann. § 30-14-2801 et seq. (Montana Consumer Data Privacy Act, as amended by SB 297, eff. Oct. 1, 2025); Colo. Rev. Stat. § 6-1-1308.5 (added by SB 24-041, eff. Oct. 1, 2025).

AI Chatbot Laws Are Adding Privacy Rules, Especially for Minors

Most health AI laws focus on use, safety, accountability or transparency rather than the full data lifecycle.

AI DATA AND ADVERTISING LIMITS

Data-sharing and ad-targeting limits for chatbots.



UTAH | Utah Code §§ 13-72a-201, -202

Mental health chatbots: no sale or sharing of a Utah user's identifiable health information or input with third parties, except narrow cases.

No use of user input to decide whether, what or how to advertise; in-chat ads must be labeled, with sponsorships disclosed.



CALIFORNIA | AB 1979 (eff. Jan. 1, 2027)

Health chatbots covered by CMIA (California's mini-HIPAA), including health data users share or the chatbot infers.

AI AND MINORS: PRIVACY PROTECTIONS

Ad, data-use and parental-control rules for youth.



CALIFORNIA | Bus. & Prof. Code § 21810 et seq.

Child accounts: memory off by default (parent-only setting); no behavior- or data-targeted ads. No data sale. Age checks eff. Jan. 1, 2027; most rules July 1, 2027.



COLORADO | HB 26-1263 (eff. Jan. 1, 2027)

Operators must estimate user age; for known minors (under 18), minors and parents get tools to manage privacy and account settings.



IDAHO & NEBRASKA | SB 1297; LB 525

Privacy and account tools for account holders (NE: minors); parents manage settings for under-13s, with risk-based tools for parents of teens.

Sources: Manatt Health, Health AI Policy Tracker (July 30, 2026); Manatt Health, Q3 2026 Health AI Tracker newsletter (draft); Manatt, 2026 State AI Legislation Tracker; Cal. AB 1979 (eff. Jan. 1, 2027); Utah Code Ann. §§ 13-72a-201, -202 (eff. May 7, 2025); Cal. Bus. & Prof. Code §§ 21810–21818 (SB 1119, Stats. 2026, ch. 190); Colo. HB 26-1263 (signed May 29, 2026; eff. Jan. 1, 2027); Idaho SB 1297; Neb. LB 525 (signed Apr. 14, 2026; eff. July 1, 2027).



Emma Beavins

Marketing and Content Lead
Vega Health

VegaHealth

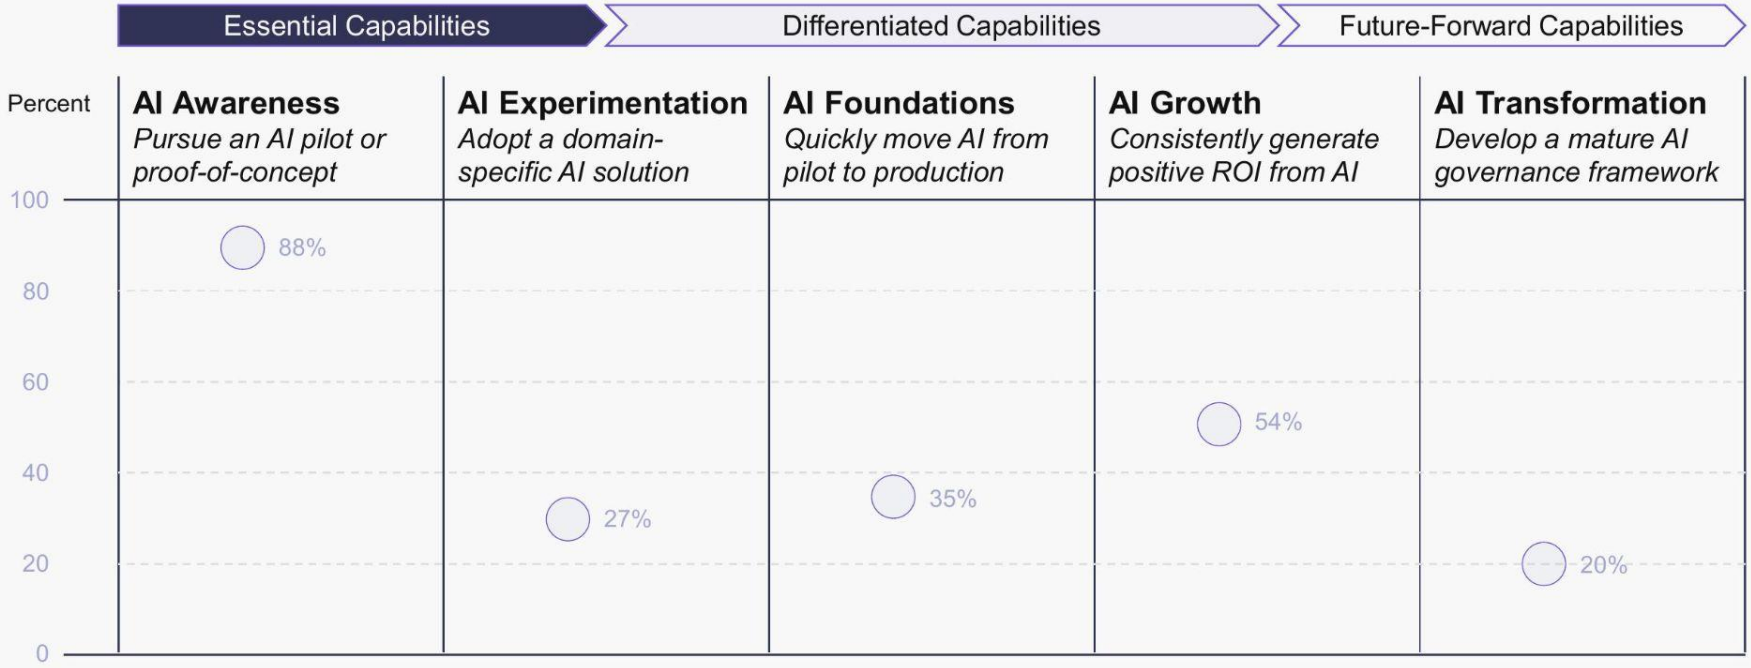
Health Data Privacy in the AI Era: The State of Health AI Adoption

October 2026

PROPRIETARY AND CONFIDENTIAL

Healthcare's AI maturity curve is lagging

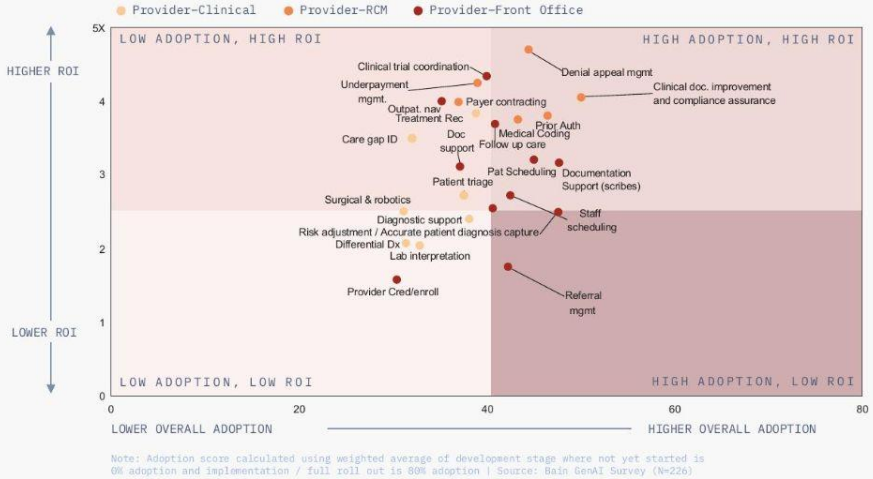
Health systems spent **$1.4 billion** on AI in 2025 – but most don't have the essential foundations to deploy AI quickly and realize value from investments



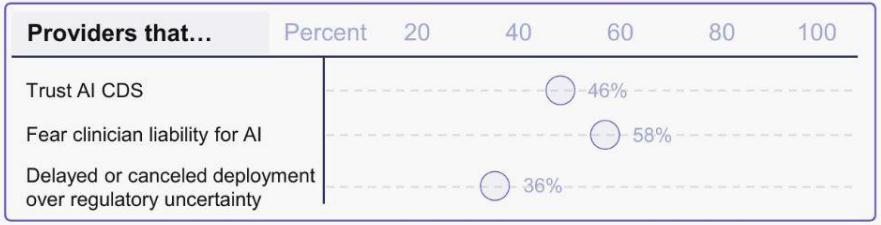
VegaHealth

Health AI adoption varies by use case and system

Bessemer Venture Partners reported that clinical AI is being adopted slower and scaling less often…



Note: Adoption score calculated using weighted average of development stage where not yet started is 0% adoption and implementation / full roll out is 80% adoption | Source: Bain GenAI Survey (N=226)



…But there is a long history of predictive AI/ML models deployed in healthcare settings, including:

01 Emergency department triage and length of stay model

02 Deterioration and sepsis early prediction bundle

03 Readmission risk and mortality prediction models

04 Postpartum depression predictive model

Source: Bessemer Venture Partners, "The 2026 Healthcare AI ROI Scorecard."

Health systems must navigate data access, privacy, and security across the AI portfolio

AI models typically access data via FHIR, HL7, or data lake integrations – examples of input data elements include:

- Patients & Demographics
- Encounters & Diagnoses
- Procedures
- Labs, Vitals & Orders
- Unstructured Clinical Notes
- SDOH Data & Flowcharts

AI solutions transform data inputs into different outputs depending on model type (predictive or generative)

Input: Model features are mapped to health system data elements

Patient Cohort: Models then run on specific patient cohorts – this varies by use case (e.g., the entire inpatient population, individually-chosen patients, etc)

Output: Models produce different outputs (scores, thresholds, free text summaries or recommendations, etc)

Privacy
Health systems – and their AI vendors – must comply with a range of state and federal data privacy laws

Security
Health systems must manage AI vendor cybersecurity risk – just like other software or device vendors

Governance
Health systems need to know whether AI models operate safely and effectively across the AI product lifecycle

VegaHealth

Bridging the gap between go-live and sustained outcomes requires lifecycle optimization

Problem Identification & Definition. Scope specific clinical problems with well-defined champions and clear ROI that a potential AI solution is expected to deliver

Repeatable Procurement Process. Understand where to buy or build – and develop repeatable processes to reduce the cost of each point solution implementation (and the cost of failed pilots/ongoing configurations)

Development & Design



Align on Success Metrics
Understand what clinical or financial outcomes represent a successful implementation



Design Expected Workflow
Understand how a new AI solution will fit into end-user workflows – and its purpose

Pre-Deployment Evaluation



Conduct Evaluation
Know how a model will perform on local patient data – particularly for safety net



Mitigate Risk
Determine whether performance is acceptable and mitigate known risks

Implementation & Change Management



End-User Education
Build AI literacy and prepare end-users to leverage the solution responsibly



Change Management
Establish guardrails around use of an AI solution to optimize implementation

Lifecycle Management & Governance



Establish Governance
Understand where. AI is used, log inputs and outputs, and track performance

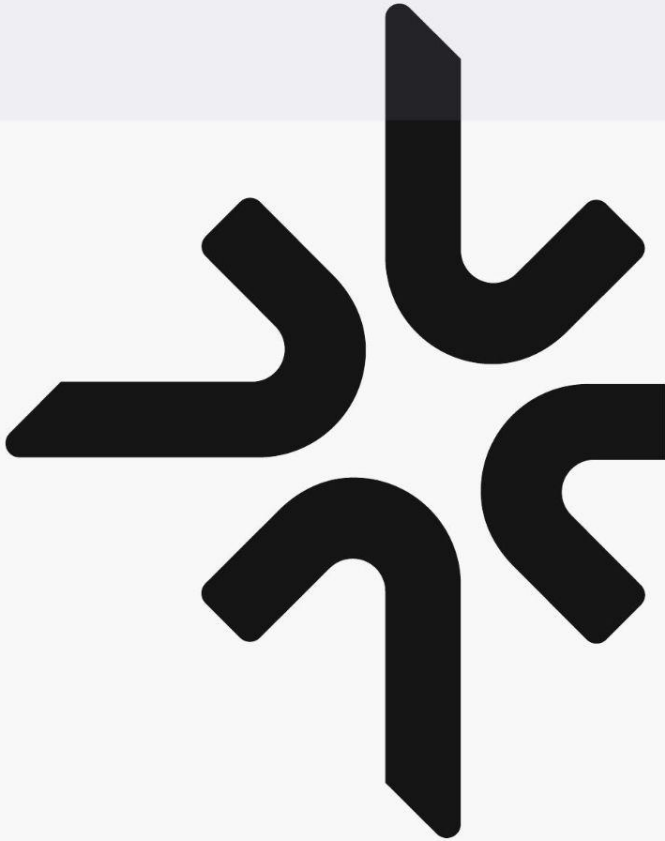


Monitor Outcomes
Track AI performance to determine whether use should continue

Contact

Emma Beavins
Content and Marketing Lead at Vega Health

Emma.Beavins@vegahealth.com





Grace Cordovano, Ph.D., BCPA

Founder, Enlightening Results
Co-Founder, Unblock Health

Health AI and the Patient Experience: Privacy, Consent, & Trust

Grace Cordovano, PhD, BCPA

Enlightening Results | October 22, 2026

Patients, Carepartners, and Caregivers Use AI

AI is already part of the patient experience and the individual and family health journey.

Patients' and families' willingness to use AI coexists with concerns about accuracy, privacy, and control.

Individuals turn to AI for a broad spectrum of reasons when it comes to health.

Patients may share sensitive health information while having concerns about privacy, especially when they feel they have few or no alternatives.

Using AI does not tell us what patients, carepartners, and caregivers understand, accept, or can control about how their information is used.

How Patients, Carepartners, and Caregivers Use AI

- Patients, carepartners, and caregivers use AI to:
  - Understand medical records, test results, peer reviewed publications
  - Prepare appointment questions
  - Explore and strategize treatment options, weighing outcomes & risks
  - Organize complex information
  - Draft insurance or administrative correspondence
- For the first time, patients have access to their own personal support team: AI thought partner, personal assistant, researcher, educator, translator, navigator, and orchestrator.
- Some patients may feel more comfortable asking sensitive health questions privately through AI. That feeling of privacy does not establish how their information is stored, accessed, or reused.

Privacy and Safety Concerns Arise Across the AI Patient, Carepartner, Caregiver Experience

Inflection Point	Questions
Choosing & sharing	What information are we sharing or connecting? Who receives it, and what protections and choices do we have, where we are?
Recording & collecting	Is AI capturing our visit or processing our records? What information is collected, and will we be told? How will the information be used?
Storing & reusing	Who can access our information? How long is it kept, could it be shared, used for training, used to identify us, or harm us?

Inflection Point	Questions
Interpreting & predicting	Does the information accurately reflect our experience? What summaries, risk scores, or predictions are generated, how do they impact our care? Can we see them?
Influencing care & coverage	Could an AI output affect treatment, prior authorization, coverage, or payment? Who reviews its use?
Correcting & challenging	Can we dispute inaccurate information, request corrections, challenge a decision, and obtain meaningful human review?

What Meaningful Consent & Transparency *Should* Look Like

- Patients, carepartners, and caregivers want to be informed when AI is used to make decisions about care, coordination of care, or is used for communications.
- What needs to be explained before use?
  - What is AI doing?
  - What information is collected, shared, retained, or reused?
  - What choices and alternatives are available?
  - What happens if I decline or change my mind later?
  - How does AI impact my costs of care?
  - Is opting-out or deletion available, what limits apply, and what happens to information already shared?
  - Was this tool/technology developed with patients like me?
  - What if harm happens? Who do I notify? What if success happens? Who do I contact to celebrate a great outcome because of AI?
- Continue transparency after use: provide access to relevant outputs and practical ways to raise concerns, request corrections, and seek human review.

Balancing Patient Autonomy, Access, Privacy, and Trust

- People should not have to trade privacy for help understanding their health or to access care.
- Patients and families deserve privacy-respecting, transparent tools that are understandable, accessible, and affordable without placing the entire burden of evaluating risk and privacy protections on the patient.
- Do privacy-enhancing approaches affect accuracy, accessibility, fairness, harms, and trust across different patient populations?
- To build trust, patients, caregivers, and carepartners should help define AI governance longitudinally, especially as technology, uses, and risks evolve.

Contact

Grace Cordovano, PhD, BCPA

Board Certified Patient Advocate

enlighteningresults@gmail.com