California Attorney General's Legal Advisory on the Application of Existing California Law to Artificial Intelligence in Healthcare

The California Attorney General's Office (AGO) issues this advisory to provide guidance to healthcare providers, insurers, vendors, investors, and other healthcare entities that develop, sell, and use artificial intelligence (AI) and other automated decision systems[1] about their obligations under California law, including under the state's consumer protection, civil rights, competition, and data privacy laws.[2]

Artificial Intelligence in the Healthcare Sector

AI systems are already widespread within healthcare. As of May 2024, the federal Food and Drug Administration (FDA) had authorized for medical use 981 artificial intelligence or machine learning software devices, and counting.[3] These and other AI systems are being used to guide medical diagnosis and treatment decisions. Hospitals and insurers routinely use non-FDA-approved AI systems for tasks such as appointment scheduling, medical risk assessment, and bill processing.

AI tools have the potential to help improve patient and population health, increase health equity, reduce administrative burdens, and facilitate appropriate information sharing. At the same time, AI risks causing discrimination, denials of needed care and other misallocations of healthcare resources, and interference with patient autonomy and privacy. For example, AI models trained on data that reflect existing biases in healthcare delivery can exacerbate health inequity.[4] Many patients are not aware of when and how AI systems are used in connection with their healthcare. Moreover, AI systems are novel and complex. Their inner workings are often not understood by the healthcare providers using AI, let alone patients receiving care.

Healthcare-related entities that develop, sell, or use AI systems must ensure that their systems comply with laws protecting consumers. This requires understanding how AI systems are trained, what information the systems consider, and how the systems generate output. Developers, researchers, providers, insurers, and related organizations should ensure that AI systems are tested, validated, and audited to ensure that their use is safe, ethical, and lawful, and reduces, rather than replicates or exaggerates, human error and biases. They should also be transparent with patients about whether patient information is being used to train AI and how providers are using AI to make decisions affecting health and healthcare.

---

1   While the definition of AI may vary depending upon the context, for the purposes of this advisory, AI includes "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to—(A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action." (15 U.S.C. § 9401(3).) California has also recently passed a law defining the term in certain instances as "an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments." (See Gov. Code § 11546.45.5 et seq., added by AB 2885, Stats. 2024, ch. 843.) "Automated decision systems" means "a computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that is used to assist or replace human discretionary decisionmaking and materially impacts natural persons." (Gov. Code, § 11546.45.5, subd. (a)(2).)

2   This advisory provides the AGO's guidance on general application of California law to AI and other automated decision systems. This advisory does not address all potential violations or avenues of enforcement for the identified laws, nor does it identify all laws that may apply to AI.

3   U.S. Food & Drug Administration, Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices (last updated May 13, 2024), https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices.

4   See, e.g., Jonathan Huang et al., Journal Evaluation and Mitigation of Racial Bias in Clinical Machine Learning Models: Scoping Review, JMIR Medical Informatics, 10(5), E36388 (May 2022), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9198828/.

For example, it may be unlawful in California to:

- Deny health insurance claims using AI or other automated decisionmaking systems in a manner that overrides doctors' views about necessary treatment.

- Use generative AI or other automated decisionmaking tools to draft patient notes, communications, or medical orders that include erroneous or misleading information, including information based on stereotypes relating to race or other protected classifications.

- Determine patient access to healthcare using AI or other automated decisionmaking systems that make predictions based on patients' past healthcare claims data, resulting in disadvantaged patients or groups that have a history of lack of access to healthcare being denied services on that basis while patients/groups with robust past access being provided enhanced services.

- Double-book a patient's appointment, or create other administrative barriers, because AI or other automated decisionmaking systems predict that patient is the "type of person" more likely to miss an appointment.

- Conduct cost/benefit analysis of medical treatments for patients with disabilities using AI or other automated decisionmaking systems that are based on stereotypes that undervalue the lives of people with disabilities.

The AGO recognizes that the California Legislature and regulatory agencies continue to develop laws and regulations addressing emerging technology. This advisory provides guidance on the application of existing California law to AI use in healthcare. This advisory does not encompass all possible laws that may apply to health AI, including applicable federal requirements, such as the FDA's regulation of software as a medical device and research into AI in medicine; the Federal Trade Commission Act; the U.S. Department of Health and Human Services (HHS) Assistant Secretary for Technology Policy and Office of the National Coordinator for Health Information Technology standards and final rule applying Section 1557 (the Affordable Care Act's non-discrimination mandate) to automated patient care decision support tools, and its guidance to Medicare Advantage plans on use of AI and other forms of automated decisionmaking; the National Institute of Standards and Technology's draft AI risk management framework; and the Biden administration Executive Order on AI and draft guidelines of the Office of Management and Budget on AI.

CONSUMER PROTECTION, CIVIL RIGHTS, COMPETITION, AND PATIENT PRIVACY LAWS PROVIDE BROAD PROTECTIONS FOR CALIFORNIANS

A. California's Health Consumer Protection Laws

California's Unfair Competition Law protects the state's residents against unlawful, unfair, or fraudulent business acts or practices, including business practices used in the practice of medicine. (Bus. & Prof. Code, § 17200 et seq.) The law was intentionally written with "broad, sweeping language" to protect Californians from obvious and familiar forms of fraud and deception as well as new, creative, and cutting-edge forms of misleading behavior. (*People ex rel. Mosk v. Nat'l Research Co.* (1962) 201 Cal.App.2d 765, 772.) In addition, a violation of any other state, federal, or local law is "independently actionable" under the Unfair Competition Law. (*Farmers Insurance Exchange v. Superior Court* (1994) 2 Cal.4th 377, 383.) Thus, the scope of the Unfair Competition Law incorporates numerous laws that may apply to AI in a variety of contexts, such as the protections against false advertising and anticompetitive practices described in Attorney General Bonta's recent general consumer legal advisory on AI.[5]

Practices that deceive or harm consumers fall squarely within the purview of the Unfair Competition Law, and traditional consumer legal protections apply equally in the AI context. This includes creation, marketing, or dissemination of an AI system that does not comply with civil rights, privacy, false advertising, competition, and other laws. State law additionally prohibits payment of referral fees or kick-backs for medical services and other types of fraudulent billing, such as use of AI to generate fraudulent bills or inaccurate upcodes of patient records. (Health & Saf. Code, § 445; Welf. & Inst. Code, §§ 14107, 14107.2.) Businesses may also be liable for supplying AI tools when they know, or should have known, that AI will be used to violate the law. (See, e.g., *People v. Toomey* (1984) 157 Cal. App.3d 1, 15 [liability under section 17200 can be imposed for aiding and abetting].)

5   https://oag.ca.gov/system/files/attachments/press-docs/Legal%20Advisory%20-%20Application%20of%20Existing%20CA%20Laws%20to%20Artificial%20Intelligence.pdf

California's professional licensing laws provide additional standards to which licensed medical professionals must adhere. (Bus. & Prof. Code, Division 2 (commencing with Section 500).) Only human physicians (and other medical professionals) are licensed to practice medicine in California; California law does not allow delegation of the practice of medicine to AI. Licensed physicians may violate conflict of interest law if they or their family member have a financial interest in AI services and must disclose any financial conflict when consulting with AI organizations. (Lab. Code, § 139.3, subds. (a), (e).) Furthermore, using AI or other automated decision tools to make decisions about patients' medical treatment, or to override licensed care providers' determinations about what a patient's medical needs are, may violate California's ban on the practice of medicine by corporations and other "artificial legal entities" (Bus. & Prof. Code, § 2400 et seq.),[6] in addition to constituting an "unlawful" or "unfair" business practice under the Unfair Competition Law.

Recent amendments to the Knox-Keene Act and California Insurance Code limit health care service plans' ability to use AI or other automated decision systems to deny coverage. (See Sen. Bill No. 1120 (2023-2024).) When employed for utilization review or management purposes, a plan cannot use these types of tools to "deny, delay, or modify health care services based, in whole or in part, on medical necessity." (Health & Saf. Code, § 1367.01, subd. (k)(1); Ins. Code, § 10123.135, subd. (j)(2).) Instead, plans must ensure that AI and other software:

- Does not supplant a licensed health care provider's decisionmaking;

- Bases decisions on individual enrollees' own medical history and clinical circumstances;

- Does not discriminate, and is applied fairly and equitably;

- Is open to inspection and audit by relevant state agencies;

- Is periodically reviewed and revised to maximize accuracy and reliability;

- Does not use patient data beyond its intended and stated purpose; and

- Does not directly or indirectly cause harm to the plan enrollee.

(Health & Saf. Code, § 1367.01, subd. (k)(1)(A-K); Ins. Code, § 10123.135, subd. (j)(1)(A-K).)

B. California Anti-Discrimination Laws

California law prohibits discrimination by any entity or individual receiving "any state support," including an "entity principally engaged in the business of providing […] health care." (Gov. Code, § 11135; Cal. Code Regs., tit. 2, § 14020, subd. (m)(6)(B); see also id. at (ii) [covered programs or activities include provision of health services].) Discrimination is prohibited based on any or a combination of the following classifications: "sex, race, color, religion, ancestry, national origin, ethnic group identification, age, mental disability, physical disability, medical condition, genetic information, marital status, or sexual orientation." (Gov. Code, § 11135; Cal. Code Regs., tit. 2, § 14000, subd. (e).)

This non-discrimination mandate covers healthcare programs or activities broadly because "state support" may come in the form of "any payments, subsidies, or other assistance extended to any person, agency or entity providing insurance, *including health-related insurance coverage for payments to or on behalf of a person obtaining health-related insurance coverage from that entity* […]." (*Id*. § 14020, subd. (ww)(5) (emphasis added).). For example, this includes state Medi-Cal services. And the non-discrimination mandate extends to all "operations of the covered entity […] even if only one part of the covered entity receives state support," including "any service, activity, financial aid or benefit provided in, at or through a facility that is or was provided by the state or any state agency or with the aid or benefit of state support or other funds or resources." (*Id*. § 14020, subd. (ii)(1-2).)

---

6   See California State Library, The Corporate Practice of Medicine in a Changing Healthcare Environment (April 2016), https://www.library.ca.gov/wp-content/uploads/crb-reports/CRB_CPM_Final.pdf, [noting that California "maintains one of the most comprehensive bans [on corporate practice of medicine] in the country."].)

These rules prohibit the types of discriminatory practices likely to be caused by AI, including disparate impact discrimination (also known as "discriminatory effect" or "adverse impact") and denial of full and equal access.[7] (Cal. Code Regs., tit. 2, § 14027, subd. (b)(3).) For example, an AI system that makes less accurate predictions about demographic groups of people who have historically faced barriers to healthcare (and whose information may be underrepresented in large datasets), though facially neutral, may have a disproportionate negative impact on members of protected groups.[8] Classifications that are protected under section 11135 may frequently overlap with lower income and social marginalization. Even if such models are applied to all patients regardless of race, they may still cause disparate impact discrimination because "identical treatment may be discriminatory." (*Id*. § 14025, subd. (a)(3).) A disparate impact is permissible only if the covered entity can show that the AI system's use is necessary for achieving a compelling, legitimate, and nondiscriminatory purpose, and supported by evidence that is not hypothetical or speculative. (*Id*. § 14029, subd. (c)(1, 2).)

Although a policy or tool may be facially neutral, healthcare entities may not simply ignore or avoid data regarding inequity relating to race, gender, or another protected classification. To the contrary, recipients of state support may be required or permitted to take ameliorative steps to overcome the effects of past discrimination, or prevent new discrimination.[9] (*Id*. § 14053; *see also id*. at § 14003, subd. (b) (California regulations should not be interpreted to adversely impact programs or activities that benefit protected subgroups in order to overcome effects of past exclusion or reduced access).

Unfortunately, real-world examples of AI healthcare systems incorporating societal and other biases into their decision making already exist.[10] Indeed, the AGO is investigating potential discrimination by AI algorithms and other automated decisionmaking products used by California healthcare entities. Developers, vendors, and users should take proactive steps when designing, acquiring, and implementing health AI to ensure that these systems do not have a discriminatory impact.

The use of AI in healthcare is subject to additional state laws prohibiting discrimination against healthcare consumers in various settings, such as:

- California's Unruh Civil Rights Act, which prohibits arbitrary and intentional discrimination by businesses, including those providing healthcare services. (Civ. Code, § 51, subd. (b); Ins. Code § 1861.03 (applying Unruh Act to insurance).

- The rights of people with disabilities to access healthcare, which are protected through additional specific disability rights statutes. For more details, see Legal Rights of Persons with Disabilities: Access to Healthcare for People with Disabilities.

- California's Insurance Code, which prohibits discrimination regarding ratemaking, claims handling, and reviewing insurance applications. For more details, see the California Insurance Commissioner's Bulletin 2022-5, Allegations of Racial Bias and Unfair Discrimination in Marketing, Rating, Underwriting, and Claims Practices by the Insurance Industry.

---

7   See, e.g., Caleb J. Colón-Rodríguez, Shedding Light on Healthcare Algorithmic and Artificial Intelligence Bias, U.S. Department of Health and Human Services, Office of Minority Health (July 12, 2023), https://minorityhealth.hhs.gov/news/shedding-light-healthcare-algorithmic-and-artificial-intelligence-bias.

8   See, e.g, Ziad Obermeyer et al., Dissecting Racial Bias in an Algorithm used to Manage the Health of Populations Science 336, 447-453 (Oct. 25, 2019) (finding that a model intended to predict patient risk disproportionately underestimated the needs of Black patients in comparison to white patients).

9   Generally speaking, California state law and policy supports actions taken to eliminate existing health disparities and inequities, including for vulnerable communities who have experienced socioeconomic disadvantage and historical injustice. (See, e.g., Health & Saf. Code, § 131019.5.) These include interventions "to address unequal starting points and drive equal outcomes so all Californians may reach their full potential and lead healthy and rewarding lives." (Governor's Exec. Order No. N-16-22 (Sept. 13, 2022).) The State of California's 2023 Reparations Report, which focused on the impact of African American slavery and its ongoing legacy of disparate harms to health, contains extensive discussion of the need for proactive measures to promote health equity. (See California Task Force to Study and Develop Reparations Proposals for African Americans, Final Report (June 29, 2023), https://oag.ca.gov/system/files/media/full-ca-reparations.pdf.)

10   See, e.g., Eliane Roosli et al., Bias at Warp Speed: How AI May Contribute to the Disparities Gap in the Time of COVID-19, 28(1) Journal of the American Medical Informatics Association, 190-192 (2021); Judy Wawira Gichoya et al., AI Recognition of Patient Race in Medical Imaging: a Modelling Study, 4 The Lancet 6, E406-E414 (May 11, 2022).

- California's Health and Safety Code requirement that licensed California hospitals have a policy of non-discrimination in access to emergency healthcare services. (Health & Saf. Code, § 1317.3, subd. (b).)

- The California Fair Employment and Housing Act (FEHA) also protects Californians from harassment or discrimination in healthcare employment, including discrimination carried out or facilitated by AI. (Gov. Code, § 12900 et seq.)

C. California's Patient Privacy and Autonomy Laws

Vast quantities of patient data underlie the massive growth in the health AI sector. Data is used to build and train AI and to render decisions that impact health services. Developers and entities that use AI in healthcare should carefully monitor training data, inputs, and outputs to ensure respect for Californians' rights to medical privacy.

California state medical privacy laws provide protections that are, in some cases, more stringent than federal health privacy laws like HIPAA (the Health Insurance Portability and Accountability Act of 1996, 45 C.F.R. Parts 160 and 164).[11] The Confidentiality of Medical Information Act (CMIA) and the Information Practices Act govern use and disclosure of Californians' medical information. Covered entities must preserve confidentiality of patients' medical information and ensure that patients have access to that information. (Civ. Code, §§ 56.10, 56.26, 1798.25.) Sensitive information, including mental and behavioral healthcare and reproductive and sexual healthcare (e.g., abortion and gender affirming care), receive heightened protections. (Civ. Code, § 56.05, subd. (s).) Medical privacy laws apply to governmental healthcare agencies,[12] medical providers, and insurance plans, as well as businesses that offer software or hardware to consumers for the purposes of managing medical information, diagnosis or treatment, or management of medical conditions, via mobile applications or other related devices. (Civ. Code, § 56.06, subds. (a), (b).)

California law requires that physicians provide information that a reasonable person in the patient's position would need for informed consent to a proposed course of treatment. (Cal. Code Regs., tit. 9, § 784.29, subd. (a) [patients' rights in mental health rehabilitation centers], tit. 22, § 70707 [patient rights in acute care hospitals].) Providers should consider whether this applies to their use of AI tools, as a majority of Californians are currently uncomfortable with use of AI in connection with healthcare.[13] If a patient is asked to participate in a medical experiment using AI systems, they are entitled to California's "Experimental Subject's Bill of Rights," including information explaining the procedures to be followed in the medical experiment, and drugs and devices used. (Health & Saf. Code, § 24172.)[14]

Significant recent amendments to the CMIA require that providers and electronic health records (EHR) and digital health companies enable patients to keep their reproductive and sexual health information confidential and separate from the rest of their medical records.[15] (Civ. Code, § 56.101, subds. (a), (c).) They must prevent disclosure, access, transfer, or processing of this information to individuals and entities outside of California. (*Id*. subd. (c)(1)(D).) As developers and users of EHRs and related applications increasingly incorporate AI, they must ensure compliance with CMIA and limit access and improper use to sensitive information.

The CMIA also imposes independent requirements on healthcare providers, insurers, and others to get patients' consent before disclosing medical information. (Civ. Code, § 56.10, subd. (a).) The Genetic Privacy Information Act provides special protections for individuals' genetic data, and California healthcare service plans and other entities are prohibited from disclosing to third parties the results of genetic tests without the patient's permission. (Civ. Code, §§ 56.17, 56.18, et seq.) "Dark patterns"—user interfaces "designed or manipulated with the substantial effect of subverting or impairing user autonomy, decisionmaking, or choice," including those generated by AI—cannot be used to obtain patient consent. (Civ. Code, § 56.18, subd. (b)(6).) Under the Patient Access to Health Records Act, California patients and their representatives have the right to obtain their own medical records. (Health & Saf. Code, §§ 123110, et seq.) Likewise, the Insurance Information and Privacy Protection Act gives healthcare consumers the rights to determine what information has been collected about them, and the reasons for adverse decisions. (Ins.

---

11 For general information about California patient privacy rights, see the AGO's Consumer Guide to Health Information Privacy in California.

12 For more information about state and local governmental agency patient privacy obligations, see the California Center for Data Insights and Innovation's Statewide Health Information Policy Manual.

13 California Health Care Foundation (CHCF), The 2024 CHCF California Health Policy Survey (Jan. 2024), https://www.chcf.org/wp-content/uploads/2024/01/2024CHCFCAHealthPolicySurvey.pdf.

14 For more information, see the AGO's resources on California Informed Consent Form Guidelines.

15 For more information about California's legal protections of reproductive healthcare, see the AGO's resources on Reproductive Rights.

Code, § 791.) Developers and users of AI must have sufficient control over their systems to ensure that Californian patients' rights to privacy and autonomy are not compromised.

Apart from these healthcare-specific privacy laws, California has general privacy laws that apply to the use of AI. For information concerning California state privacy laws and AI, including the constitutional right to privacy that applies to both government and private entities (see *Hill v. National Collegiate Athletic Assn.* (1994) 7 Cal.4th 1, 20) and the California Consumer Privacy Act, see Attorney General Bonta's recent general consumer legal advisory on AI.

HEALTHCARE ENTITIES SHOULD REMAIN VIGILANT ABOUT OTHER LAWS AND REGULATIONS WHICH MAY BE APPLICABLE TO AI TECHNOLOGIES

Beyond the laws and regulations discussed in this advisory, other California laws—including tort, public health, charitable trusts, competition, and criminal laws—apply equally to AI systems as they do to non-AI system. Put another way, conduct that is illegal without the involvement of AI is equally unlawful if AI is involved, and the fact that AI is involved is not a defense to liability under any law.

This overview is not intended to be exhaustive. Laws and regulations will undoubtedly continue to evolve in the face of new technology. But healthcare entities that develop or use AI should not wait to ensure that they comply with all state, federal, and local laws that may apply to their use of AI. That is particularly so when AI is used or developed for applications that carry a potential risk of harm to patients, healthcare systems, or the public health writ large.