**Ten Principles for Protecting Consumers in AI-Enabled Healthcare**

*Human Oversight and Transparency*

1. **AI should support—but never replace—clinical judgment**. AI should be used to support healthcare professionals and improve patient care, not replace their clinical judgment or role in making decisions about a patient's care. AI-generated recommendations should not be the sole basis for denying, delaying, or terminating medically necessary care. In addition, when AI is used in insurance decisions that could deny or delay access to healthcare, a licensed healthcare professional must make the final determination, exercise meaningful review of the patient's individual circumstances, and possess independent authority to overturn any AI recommendation.
2. **Consumers should receive clear disclosure when AI plays a significant role in their healthcare**. Consumers should receive clear, understandable disclosure when they are interacting directly with AI (e.g., through a chatbot) or when AI materially influences decisions about their care or insurance coverage. Transparency should be meaningful and proportionate to the role AI plays.
3. **Consumers should have the right to understand and challenge AI-assisted healthcare decisions**. Consumers should retain robust rights to challenge AI-assisted decisions, obtain human review, understand why an adverse decision was made, correct inaccurate information underlying it, and access appeal mechanisms without administrative burden, including clear deadlines for insurer and provider responses.
4. **Consumers should be able to understand the capabilities and limitations of AI health tools**. AI tools that provide health information or recommendations should clearly communicate what they can and cannot reliably do, including when consumers should seek professional medical care. They shouldn't present uncertain or unvalidated information as authoritative medical advice. In addition, when relevant, AI health tools should clearly communicate whether they are FDA approved/cleared/authorized and for which uses.

*Safety, Effectiveness, and Equity*

5. **AI must be safe, accurate, and reliable for its intended use**. Health AI should be validated for its intended use before deployment and monitored afterward for accuracy, reliability, safety, and changing performance.
6. **AI must work fairly for all patients**. Training data and evaluation should adequately represent the populations the technology will serve. AI systems should be tested for bias and disparate performance across patient groups, including by race, ethnicity, sex, age, disability, language, and socioeconomic status. Testing should occur before and after deployment, with corrective action when disparities emerge. Because bias and algorithmic drift (long-term change in an algorithm's outputs) often emerge well after

initial deployment, AI systems should undergo mandatory, independent audits for disparate impact on a regular schedule.

*Privacy and Data Protection*

7. **Consumers' health information should be protected and secure wherever it is collected**. Sensitive health information should be protected regardless of where it is collected. Consumer privacy protections comparable to those under HIPAA should apply to health information collected outside of HIPAA-covered settings, including for wearable health devices, apps, and related websites, with appropriate safeguards governing how that information is collected, used, stored, and shared. In addition, organizations that collect, use, store, or share health information should maintain strong security safeguards to protect that information from unauthorized access, breaches, theft, and misuse.
8. **Companies should collect and use only the health data necessary to serve consumers**. AI systems can collect, analyze, and combine enormous amounts of personal information, including highly sensitive health data. AI developers and healthcare organizations should collect and retain only the information necessary to provide the service a consumer requests and limit its use and sharing to that purpose. Sensitive health information should not be sold to advertisers or repurposed for unrelated uses, such as targeted advertising or training AI models.
9. **Consumers should have meaningful control over their health data—including the right to correct and delete their information, and, where appropriate, to be forgotten**. Consumers must be able to access the personal information an AI system holds about them, easily correct inaccurate information, and delete information they no longer want retained.

*Accountability and Enforcement*

10. **Healthcare providers, insurers, AI developers, and other entities should have clear oversight responsibilities and be accountable when AI causes harm**. Entities that deploy AI must maintain public reporting of safety incidents (e.g., to appropriate regulatory authorities) and be prohibited from using pre-dispute mandatory arbitration clauses or liability waivers to block consumer claims. Organizations should provide a clear point of contact for consumers' questions and requests, including a straightforward mechanism to report suspected errors and receive a timely response. Finally, Congress should also establish a private right of action with statutory penalties for when health AI causes harm to consumers.