

This PDF is available at <https://nap.nationalacademies.org/29562>



Exploring the Application of Artificial Intelligence in Obesity Care: Proceedings of a Workshop—in Brief (2026)

DETAILS

19 pages | 8.5 x 11 | PDF

ISBN 978-0-309-74972-5 | DOI 10.17226/29562

CONTRIBUTORS

Erin Hammers Forstag, Rapporteur; Roundtable on Obesity Solutions; Food, Nutrition, and Agriculture Program Area; Center for Health, People, and Places; National Academies of Sciences, Engineering, and Medicine

SUGGESTED CITATION


National Academies of Sciences, Engineering, and Medicine. 2026. *Exploring the Application of Artificial Intelligence in Obesity Care: Proceedings of a Workshop—in Brief*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/29562>.

BUY THIS BOOK

FIND RELATED TITLES

Visit the National Academies Press at nap.edu and login or register to get:

- Access to free PDF downloads of thousands of publications
- 10% off the price of print publications
- Email or social media notifications of new titles related to your interests
- Special offers and discounts



All downloadable National Academies titles are free to be used for personal and/or non-commercial academic use. Users may also freely post links to our titles on this website; non-commercial academic users are encouraged to link to the version on this website rather than distribute a downloaded PDF to ensure that all users are accessing the latest authoritative version of the work. All other uses require written permission. ([Request Permission](#))

This PDF is protected by copyright and owned by the National Academy of Sciences; unless otherwise indicated, the National Academy of Sciences retains copyright to all materials in this PDF with all rights reserved.

Convened June 3–4, 2026

Exploring the Application of Artificial Intelligence in Obesity Care

On June 3 and 4, 2026, the Roundtable on Obesity Solutions at the National Academies of Sciences, Engineering, and Medicine hosted the workshop Exploring the Application of Artificial Intelligence in Obesity Care. Bruce Y. Lee, chair of the workshop planning committee, welcomed participants and reviewed the workshop structure. The workshop featured expert presentations and panel discussions on a variety of topics, including two conversations to demonstrate what can possibly go wrong with artificial intelligence (AI) (Box 1) and what happens when someone asks AI about obesity care advice (Box 2). It began with foundational information about the basics of AI and the ethical considerations associated with its use, followed by a discussion of opportunities for the application of AI to obesity care. Subsequent sessions covered how AI could be used for the prevention, treatment, and management of obesity. The workshop closed with a discussion around the use of AI for policy making and a final presentation that captured key themes and identified next potential steps.

LANDSCAPE OF AI

Lee kicked off the workshop by defining AI, describing differences among AI tools, and identifying some of the opportunities and issues associated with AI. AI is an umbrella term that “encompasses any computer-aided technique that can replicate something that a human brain would normally do beyond simply following step-by-step directions,” said

Lee. This includes replicating human functions such as learning, reasoning, problem-solving, perception, language understanding, creating and designing. Examples of AI approaches, methods, and terminology include the following:

- Generative AI: creates content based on patterns from training data
- Natural language processing (NLP): understands, interprets, and generates human language
- Large language models (LLMs): trained on large amounts of text to perform NLP tasks
- Machine learning (ML): learns from data and makes decisions without being provided hard-coded rules
- Deep learning: subset of ML that uses multi-layered neural networks
- Agentic AI: autonomous system that sets goals, creates plans, and executes tasks with minimal or no supervision
- Foundation models: trained on very large datasets for broad application
- Multi-modal model: can process, interpret, and generate many data types simultaneously
- Artificial general intelligence (AGI): hypothetical AI form that matches or surpasses all human functions (also known as “the singularity”)

In general, AI methods and approaches fall at or between two ends of the spectrum, said Lee, including top-down, data-driven approaches that use data to identify and discover

patterns and associations, as well as bottom-up, mechanistic approaches that use particular mechanisms of cognition. Many AI tools are hybrids of these two approaches.

Training AI systems can be supervised or unsupervised, said Lee. With supervised learning, the user gives the AI algorithm/model data and what specifically the AI algorithm/model should determine from the data. With unsupervised learning, the user gives the AI algorithm/model data without the “correct answers.” And the AI algorithm/model then determines on its own what the answers should be. Semi-supervised learning is a hybrid of the two. Lee noted the similarities to how humans learn, with supervised learning in classrooms and unsupervised learning in the “real world.”

Like all technologies, said Lee, AI is neither inherently good nor bad. If AI is used to facilitate growth and identify new ways of addressing issues, it can be positive; if it is used to worsen existing problems (e.g. bias), replace humans where humans are needed, create havoc, and facilitate bad behavior, it can be negative. Lee summarized these potential risks and opportunities across different domains (Table 1-1).

In closing, Lee re-emphasized that “AI” is not just one tool or definition—it is an overarching term that represents a number of computational approaches. Over the next several years, he said, the core definition of AI will remain the same, but the specific methods, algorithms, models, and tools will evolve.

ETHICS AND GOVERNANCE

These sessions discussed the ethical considerations associated with AI in general, obesity care in particular, and the governance structures that could support responsible and ethical use.

Ethical Frameworks for AI

Scientific progress tends to outpace norms and governance, said Debra Mathews, Johns Hopkins University Berman Institute of Bioethics and School of Medicine. In the current era, technologies typically don’t remain in traditional regulatory, market, or academic “buckets,” and existing laws and regulatory structures may no longer be fit-for-purpose. The current regulatory lens tends to focus on individual benefits and harms, with no mandate or capacity to attend to broader ethical principles. According to Mathews, AI presents

many opportunities in healthcare—such as enhancing wellbeing, reducing bias and misinformation, and improving access to care, diagnosis, and treatment—but also introduces unique challenges. Mathews said that along with the familiar ethical challenges such as privacy, informed consent, and data sharing, there is a unique constellation of AI challenges. Core challenges include AI’s scale, pervasiveness, and invisibility; the ways in which bias can get “baked” into algorithms; and a lack of information about what is happening within the “black box” when an AI tool is reasoning and making decisions. These challenges, she said, add to the contextual challenges of the modern age. People live in an age of “surveillance capitalism,” where the data of individuals are collected and used to sell things to them (Zuboff, 2019). Technologies are being used to do things faster, cheaper, and at scale, and there is a mismatch between protections and risks. New technologies affect societies in uneven ways—some people benefit while some are harmed, often in systematic ways—and exacerbate existing inequities.

As AI has grown and evolved over the last decade, said Mathews, several groups have attempted to develop ethics and governance frameworks. These frameworks differ, but they tend to share five common themes: transparency, justice and fairness, non-maleficence, responsibility, and privacy (Jobin et al., 2019). In 2022, the White House’s Office of Science and Technology Policy (OSTP) published a “Blueprint for an AI Bill of Rights” (OSTP, 2022). This report involved extensive public engagement, said Mathews, and articulated five guiding principles: (1) AI systems should be safe and effective; (2) AI systems should not be discriminatory; (3) Individuals’ data should be protected; (4) Individuals should have notice and explanation of the use of AI systems; and (5) A human should be available to consider and remedy problems that occur. Mathews noted that while this blueprint is no longer national policy, it is being used at the state level to guide policy and regulatory decisions.

Americans largely agree on the need for regulation of AI systems, said Mathews. A 2025 study found that while many consumers are adopting AI tools, they have low trust in whether it will be used appropriately (Gillespie et al., 2025). Around 72 percent of U.S. consumers say more regulation is needed, and 43 percent have low confidence in commercial and government entities to develop and use AI. More Americans trust universities, research institutes, and health

Table 1-1 Risks and Opportunities Associated with AI

	Risks	Opportunities
Economy	Job losses	Job gains
	AI replacing traditional job functions and roles	Humans focus on different and new functions and roles
	Higher productivity expectations	Higher efficiency and new operations
	Major consolidation	Lower barriers to entry
Information marketplace	Misinformation	Better access to information
	Manipulation	Interactive discussion
	Increase in bias	Decrease in bias
	Overreliance on one source	Critically considering multiple sources
Legal	Help evade cybersecurity	Evaluate and improve cybersecurity
	Collect and use personal data	Track and safeguard personal data
	Intellectual property infringements	Intellectual property protections
	Loss of accountability	Help track processes
Environment	Data centers may increase pollution	Identify new ways of doing things and environmental solutions
	E-waste	
	Water depletion	
	Electricity/grid challenges	
Social governance	Colonialism	Opportunities for self-determination
	Resource exploitation	Enhanced self-governance of resources
	Models may perpetuate bias	Models may be less biased

SOURCE: Summary of comments presented by Bruce Y. Lee, June 3, 2026.

care providers to develop and use AI in the best interests of the public. In developing regulations around AI, said Mathews, there are several critical questions to consider:

- What are appropriate roles for autonomous systems?
- Who should decide?
- How do we decide?
- Whose values should guide our decision-making?
- How do we meaningfully integrate values into governance?

Ethics Throughout the AI Life Cycle

Kelly Michelson, Northwestern Feinberg School of Medicine, spoke about the ethics associated with using AI for managing obesity, emphasizing that ethical challenges arise at every stage of the AI life cycle.

Problem Selection

Before a tool is created, a decision must be made about which problem to address, said Michelson. When efforts are focused on a particular problem, resources and attention may be diverted from other areas. For example, compared to other areas, a large amount of work has focused on AI and bariatric surgery. For individuals or populations for whom surgery is not appropriate, said Michelson, this decision to focus on bariatric surgery means that there is less work happening in other areas of obesity care.

Data Creation and Acquisition

All AI technologies start with data, said Michelson. Where these data come from can make a huge difference in how technologies are used, who they serve, and the outcomes they produce. Bias is a major ethical challenge in data creation and acquisition. Systematic errors can become embedded in the data, with consequences for accuracy and equity, said Michelson. For example, some populations may be overrepresented or underrepresented in the data based on factors such as age, sex, race, ethnicity, healthcare utilization, and many others. Using these data to train AI models can create further inequities and ethical challenges. Michelson shared an example from an AI model that was designed to identify people at risk for needing a lot of medical care; the hope was to connect high-risk people with support systems and resources. The developers found that when presented with a white person and a Black person with identical clinical profiles, the model gave the white person a higher risk score. They determined that this issue arose because the model made

its risk determination not only on clinical characteristics but on health care utilization, and the model assumed that because white people used health care more frequently, they needed it more. Disparities in any of the characteristics used to train an AI model can result in bias in the resulting model, said Michelson.

There are also important ethical considerations around data privacy, security, and confidentiality. People share very personal information with their healthcare providers; if this information is shared with technology developers, there is a risk of erosion of trust in the healthcare system. Further, there are significant financial motivations involved in data sharing, and existing protections (e.g., Health Insurance Portability and Accountability Act [HIPAA]) only protect certain types of data and certain entities. On the other hand, said Michelson, data sharing has enormous value for furthering knowledge and improving care. The question is, she said, how can data be shared in a way that protects privacy while enabling innovation?

Development

In the process of developing an AI tool, people make numerous decisions about which variables to include and how these variables will be considered in the AI's analysis. As in other steps of the AI lifecycle, these decisions can have major implications. For example, said Michelson, an AI tool that was created to identify cancerous moles from images was found to be highly accurate. However, upon further examination, the developers found that the tool was not basing its decision on the mole itself, but on the inclusion of a scalar in the image. She explained that a scalar is used to track the size of a mole and is often included in an image when a clinician is worried that the mole may be cancerous. As AI becomes increasingly complex, addressing the concerns about the "black box" of decision-making will become more critical. "If we don't understand what the AI is making its decisions around, how can we make sure that it's giving us the answers that we need?" asked Michelson.

Evaluation

Determining whether an AI tool works or not is not clear-cut, said Michelson. There are implications and trade-offs associated with every decision in the evaluation process. For example, for a screening tool to do its job, decisions need to be made about how to prioritize sensitivity and specificity,

and where cutoff points should be set. These decisions could result in one group receiving better or worse treatment or having better or worse outcomes. A tool might “work” for one population but not for another, said Michelson, and it is critical that evaluations consider these differences.

Deployment

Putting AI tools into practice implicates a number of ethical issues, said Michelson. There are questions about consent, privacy, and liability, as well as broader issues of maintaining the clinician–patient relationship, environmental impact, and equitable access to these technologies.

Unique Considerations for Overweight and Obesity

In addition to these AI-specific challenges, there are also unique ethical issues related to overweight and obesity. There is ambiguity around defining overweight and obesity, and this ambiguity is reflected in the data that are being used to develop systems. There are social and cultural norms, as well as stigma, that might impact the data that are collected and the AI systems that are built. One particular ethical challenge, said Michelson, is that the growing prevalence of obesity means there are strong financial incentives to develop solutions. It is important to consider how financial motivations might impact the choices made in the development of AI tools directed at overweight and obesity, she said.

Strategies for Supporting Ethical and Responsible AI

Given the ethical challenges discussed, Michelson offered her suggestions for moving forward. Oversight is critical, she said. This oversight can be driven by a simplified ethical framework that considers, at each stage:

- Is it true?: Are the outputs of AI-based technologies trustworthy, which includes accuracy, re–producibility, and verifiability?
- Is it good?: Do the AI-based technologies improve health care delivery and benefit patients, families, clinicians, and/or public health?
- Is it wise?: Are the decisions that have been made based on experience and knowledge from different domains, and do they consider all options and potential consequences?

Michelson noted that policy development will also be key for the ethical management of AI; policies will need to balance responsible policy with enabling innovation. In addition,

health care providers will need education about AI and how to evaluate new tools. Education should focus deeply on instilling critical thinking skills at all stages of training, said Michelson. In all potential applications of AI, it will be important to consider what aspects of care most benefit from human judgement and human connection—and how to ensure that those aspects keep the human touch.

AI Governance

S. Craig Watkins, The University of Texas at Austin, told workshop participants about his work at IC², a research institute that focuses primarily on the application of AI in healthcare contexts. AI tools are becoming more widely integrated into health care settings and environments, he said, and he and his colleagues are interested in understanding how the clinical reasoning process is evolving—for instance, to what degree do these tools enhance clinical decision-making and agency, and to what degree is there a risk of diminishing clinicians’ ability to maintain their own agency and judgment? One area that Watkins works in is AI governance. He explained that AI governance refers to the “structures, processes, roles, and accountability mechanisms through which organizations oversee AI systems throughout their life cycle.” In other words, governance defines who is responsible, how decisions are made, and how compliance is monitored and enforced. Over the last several years, a set of guiding principles has emerged to create a framework for leveraging and enhancing the capabilities of AI systems and mitigating the risks. Responsible AI governance requires clear identification of all stakeholders, said Watkins, including health systems leaders, regulatory agencies, data scientists, AI engineers, clinicians, and patients. Each stakeholder has a unique and important role in the different governance mechanisms.

As an example, Watkins shared a framework that he and his colleagues developed to look at AI in healthcare through an ethical lens. The framework, from the report “Building a Healthier Future: Designing AI for Health Equity” (NAACP, 2025), defined three layers associated with specific stakeholders and responsibilities:

- Ethical and normative layer: Forms the moral and legal framework for health AI governance
- Organizational and governance layer: Embeds safety and responsibility principles into institutional and technical decision-making

- Operational and lifecycle layer: Creates practical tools and protocols that attend to safe and responsible AI across the life cycle.

Watkins illustrated how this framework would operate in the context of an AI chatbot delivering obesity care, noting that many decisions need to be made to build governance that ensures the chatbot is safe and clinically valid. Data stewards would need to define and enforce rules around which patient data the chatbot can access, such as weight history, glucose readings, and dietary logs, and determine which data are off-limits, how data are stored, and who can access the data. Another aspect of governance for a chatbot would be the area of clinical accountability and escalation, including a defined escalation protocol specifying the exact conditions under which the chatbot must transfer a patient to a human clinician (e.g., expressed suicidal ideation, disordered eating, dangerous glucose readings). Watkins observed that there have been tragic stories about how chatbots have failed users who have committed suicide or engaged in self-harm after their interactions; moving forward, chatbots need to be designed and engineered to identify these kinds of risks and escalate the user to a human professional. Finally, a third important area for governance is ongoing monitoring and model drift. There is a need to develop a monitoring protocol that can track chatbot performance metrics over time, including response accuracy, escalation rate, patient satisfaction, and correlation with clinical outcomes.

As AI moves into obesity care, said Watkins, it will be increasingly important that systems become more transparent and understandable; designers, clinicians, and users need to understand the mechanisms and impact of AI tools so that risks can be identified, mitigated, and addressed. AI governance is a critical component of this process.

APPLICATION OF AI TO OBESITY

These sessions explored ideas about potential applications of AI, how these would change obesity care, and what actions are may be needed to realize this vision.

Opportunities for AI in Obesity

“We are at a critical inflection point,” said Shamira Sridharan Weaver, Google. The rise of weight loss medications has fundamentally altered consumer biology, healthcare pathways, and the economics of food, and this rise

coincides with incredible growth in AI innovation. There are opportunities to use AI to address some of the challenges associated with the medications and to shift from treating obesity as a monolithic, weight-centric condition toward a precise, personalized metabolic health model. Weaver identified three specific opportunities: (1) AI can dramatically expand access to services such as lifestyle and nutrition guidance; (2) AI can provide personalized insights and guidance to individuals; and (3) AI can be available at all hours of the day. Weaver noted that people make decisions about food multiple times a day; a client can’t call a nutritionist for every decision but could access guidance through AI.

Weaver shared some of the technological innovations that will enable this type of personalized, continuous care. Both Google and Apple phones have health ecosystem programs that can connect data from various sources into one place and combine data such as weight, body composition, nutrition, and glucose data, which can be shared with caregivers or care providers. The data can also be uploaded to AI tools to develop plans and answer questions, said Weaver. Another tool for personalized obesity care is agentic AI. An individual could have multiple agents that play different roles and collaborate on care. For example, to answer a question about sleep, a data science agent would review all relevant data, surface these data to a domain expert agent, and then a health coach agent would give advice on the best actions to take. Before implementing this model, said Weaver, it would need to be validated using a framework such as SHARP, which involves a collaboration between the AI agents and human experts to ensure that the advice given is Safe, Helpful, Accurate, Relevant, and Personalized. SHARP is generally applicable across most agentic flows, she said, but each role-specific agent also needs to be checked for accuracy. For example, it is critical that the data science agent is pulling and analyzing the correct data, and that the domain expert agent is working with up-to-date, comprehensive, and clinically sound information.

These AI approaches will enable obesity care that is continuous instead of episodic, and that focuses on deep phenotypic personalization rather than simple dietary guidance, said Weaver. Achieving this vision will require leveraging and integrating multiple data sources and tools, including medical records, wearables, genomic information, and more. Individuals will be able to access information at all hours and receive guidance on meal planning and choosing foods that

are optimized for their own health and goals. This model also enables better collaboration with care teams because it provides a holistic picture of a person's health and behaviors on an ongoing basis, said Weaver. She emphasized that the success of this approach will require cross-disciplinary collaboration among stakeholders, including researchers, tech companies, and clinicians, along with an ethical framework that ensures privacy and equitable access.

Responsible Use of AI in Obesity

With global obesity rates high and rising, AI may be a critical tool for scaling up treatment and management, said Fatima Cody Stanford, Harvard University. AI-driven tools have unique advantages—they can offer immediate, interactive, personalized, and always-available assistance. Stanford said a growing body of evidence supports AI's role in addressing obesity, including the use of LLMs, machine learning, predictive AI, and computer vision and medical imaging. For example, chatbots have been shown to perform on par with human obesity care providers, and experts have rated AI-generated weight loss as clinically acceptable.

Stanford identified eight key opportunities for AI across the obesity care continuum:

- Lifestyle support: Personalized nutrition, exercise plans, and coaching
- Patient engagement: 24/7 motivational support, stigma reduction
- Clinical empowerment: Administrative tasks, EHR analytics
- Medication guidance: Reminders, side-effect monitoring
- Virtual assessment: BMI, waist-hip ratio
- Surgical guidance: Pre/post-bariatric education
- Predictive modeling: Risk stratification, obesity onset prediction
- Research assistance: Literature synthesis, study design, data analysis

For these opportunities to be realized, however, several critical challenges must be addressed, said Stanford. These challenges include issues around accuracy and reliability; transparency and black box opacity; privacy, legal, and regulatory uncertainty; algorithmic bias and cultural sensitivity; overreliance and deskilling; and the lack of evidence from randomized, controlled trials (RCTs).

Stanford suggested several guidelines for safe AI deployment, including keeping a “human in the loop,” using standardized evaluation frameworks for accuracy and clinical utility, implementing bias detection protocols, using diverse and representative training datasets, and developing population-specific models. Stanford said all stakeholders have roles to play to meet the shared goal of responsible, equitable, and evidence-based AI integration in obesity care. Researchers can prioritize conducting robust RCTs, develop evaluation frameworks, and systematically test for biases. Clinicians can use AI as decision support while maintaining clinical judgment and accountability, advocate for patient-centered AI, and educate patients on the appropriate use of AI for between-visit support. Policy makers can establish regulatory pathways for obesity AI tools, push for the creation of a centralized AI registry for performance tracking and bias monitoring, and fund equitable AI infrastructure.

Stanford emphasized that the issue of equity is “non-negotiable,” and she stressed that AI must serve all patients, not just those with digital access and high socioeconomic status. If AI is implemented responsibly, she said, “it has the potential to support and enhance existing obesity care pathways, offering scalable, personalized interventions that can empower patients and reduce health care burdens.”

AI IN THE PREVENTION OF OBESITY

In this session, speakers discussed AI-based approaches for obesity prevention.

AI Analytics to Understand the Systems Shaping Obesity

Obesity prevention requires systemic and individual strategies, said Kayla de la Haye, University of California, Los Angeles. Obesity is driven by dozens of interacting factors—including genes, behaviors, environments, and policies—that shape an individual's risk. De la Haye focused her remarks on healthy diets as one key element of child obesity prevention. For many families, she said, healthy eating is very difficult. She shared an example from her work in Los Angeles with low-income families. These families were highly motivated to make changes and to implement healthy habits, but their ability to do so decreased significantly when the COVID-19 pandemic resulted in unemployment and food insecurity. Families reported that buying fruits and vegetables was

beyond their means, and that they needed to allocate their food dollars toward cheaper foods that wouldn't go bad and that they knew their children would eat.

Despite more than a century of nutrition research and guidance in the U.S., most Americans do not follow the U.S. Dietary Guidelines for Americans. The vast majority of Americans exceed recommended limits for sugar, saturated fat, and sodium and under-consume vegetables, fruits, and other nutrient-dense food. Adherence to these guidelines is particularly challenging for people with less education and wealth, she said. Eating is an incredibly complex behavior influenced by many factors, including internal physiological and psychological cues that drive food choices (e.g., hunger, knowledge, beliefs, emotions); routine habits with high automaticity in response to cues; and health-related social needs (i.e., social determinants of health).

AI presents new opportunities to understand and modify eating “exposomes”—that is, the complex system of above-the-skin factors that shape what people eat. One area for opportunity is in better understanding the food environment in which people acquire and consume food. Many Americans don't have access to sufficient, healthy food; de la Haye said about 25 percent live in a food desert and 14 percent are food insecure (Rabbitt et al., 2025). Exposures to unhealthy food environments are particularly concentrated among people with low incomes, rural communities, and communities of color. A new project that used AI to analyze smartphone mobility data has revealed new insights about how people navigate their food environments (Bulle et al., 2024; Horn et al., 2023). The project looked at millions of Americans' cell phone data from 11 large metro areas, and found that urban residents traveled, on average, 4 miles to fast food outlets and 2 miles to grocery stores. Further, they found that as people move around in their daily routines, higher exposure to fast food outlets increases their likelihood of visiting a fast food outlet. These insights, said de la Haye, can inform food environment interventions to optimize efforts where most unhealthy food visits occur and where people are most influenced by food environments. For example, personalized place-based interventions could send messages to people when they are in an unhealthy food environment and provide them information about where healthy choices are located.

There are tremendous opportunities to use AI for innovative obesity prevention, said de la Haye. Technology can be used to account for complex exposomes by monitoring and understanding contexts, develop behavior change guidance that is tailored to a person's ecological context, and generate insights that can be used to advocate for food and nutrition security and equitable food access. De la Haye emphasized that achieving equity in obesity prevention will require breaking down barriers to adopting a healthy diet, exercise, and lifestyle; AI may give us the opportunity to do so.

Wearable Technology in Preventive Care

People have many decisions to make about their personal health habits, said Hannah Dimmick, City University of New York. They make these decisions based on information from a variety of sources, including personal values, beliefs, and intuition; community and social media; health care providers; and medical test results. This data ecosystem has limitations, she said; for example, lab test results are narrowly targeted and infrequently collected, and social media may present conflicting and incorrect information. Wearables, in theory, offer individuals the opportunity to make decisions about their health based on consistent, accurate, and integrated data, said Dimmick.

Early wearables included only basic descriptive metrics, such as step count, while newer devices include composite metrics with implicit value judgments. For example, a device might interpret multiple signals to quantify an individual's sleep quality. All of these approaches, however, simply provide the user with information to make decisions. The next generation of wearables, said Dimmick, will analyze adaptive metrics to make personalized recommendations and use reinforcement learning to nudge behavior. She shared details of promising research on using wearables for individualized care in different areas. One study looked at the relationship between daily activity and symptoms of myasthenia gravis (Dimmick et al., 2024). With twelve weeks of tracking data, researchers were able to see individual differences in response to activity and give personalized exercise recommendations. In the field of obesity, Dimmick said, researchers have used continuous glucose monitoring data to define different metabolic sub-phenotypes among individuals with diabetes (Metwally et al., 2026). This information gives clinicians the ability to provide more personalized and effective treatments, she said.

As these technologies become more widespread, Dimmick said, it is critical to consider the risks associated with their use. There are risks to the user, including negative psychological effects such as anxiety, overconfidence, and habituation. For example, users may develop anxiety over getting a “good” score, may ignore symptoms due to overconfidence in their wearable scores, or may eventually ignore system nudges toward healthy behaviors. In addition, there are risks that the data may provide misinformation and lead to wrong or missed treatment, as well as risks to privacy with targeted advertising based on personal details. Dimmick gave several suggestions for moving wearable technologies forward:

- Improve the data used, particularly by expanding training populations to a broader spectrum of individuals;
- Incorporate principles of physiology and psychology into algorithms;
- Use metrics that are meaningful, not just measurable, and give users more context about what information the model uses to make a conclusion;
- Make it easier to share data with healthcare providers; and
- Add user input to adaptive metrics.

Wearables and AI platforms have the potential to dramatically change the precision in the preventive care ecosystem, said Dimmick, by helping people better understand the signals from their bodies, keeping people aligned with their personal beliefs and values, and improving communication with health care providers.

Wearables for Predictive, Preventive, Personalized, and Participatory Care

Torey Lee, WHOOP, told workshop participants about the P4 Health framework, pioneered by Leroy Hood and further developed with Nathan Price, which is based on a shift from reactive disease treatment to “proactive, data-dense health optimization.” AI and wearable technologies provide the pathway to realizing the P4 vision at scale, said Lee. She shared examples of how wearable technology and AI can create *Predictive, Preventive, Personalized, and Participatory* obesity care. This framework is critical for obesity specifically, she said, because the obesity trajectory often starts years before a patient presents with overweight or obesity in the clinic. Lee noted that most wearable platforms were originally designed for healthy early adopters; designing for people living with obesity requires intentionally different design choices. She walked workshop participants through the four Ps.

Predictive

Wearables can help identify early signs of disease risk, even before the onset of symptoms. For obesity-related conditions, the detection window can be quite large, said Lee. For example, type 2 diabetes and cardiovascular disease usually develop over many years with subtle physiological changes that occur long before clinical diagnosis. The ability of wearables and AI to predict disease is still in development, said Lee. For example, a few existing products on the market use passively collected data from wearables to identify physiological changes consistent with respiratory illness prior to the onset of symptoms, as well as earlier screening of cardiovascular arrhythmias and hypertension. The real transformative potential in this area, said Lee, is the use of “foundation models” to identify chronic disease risk. Foundation models can learn rich representations of human physiology and behavior from large-scale multi-modal data; for obesity, this could mean combining physiological monitoring, nutritional data, activity patterns, sleep architecture, and clinical markers to enable predictions that no single data stream alone could support. Lee explained that these models can extract a rich signal out of noisy data to infer whether an individual’s trajectory is headed toward health or toward metabolic dysfunction.

Preventive

In the context of obesity, prevention is driven in large part by factors such as sleep, nutrition, physical activity, and stress management. Wearables and digital platforms are well-suited to capture these important behaviors continuously and often objectively. Further, said Lee, these technologies can help create a foundation for intervention. AI can connect dots in the data to reveal patterns; for example, a platform could help an individual understand how their sleep the previous night impacts their dietary choices the next day, or how stress levels drive eating patterns. Based on these patterns, the platform can nudge individuals toward healthy behaviors, using reinforcement learning models that adapt to what works best for a specific individual. She explained that reinforcement models can modify factors such as the tone or timing of messages, measure users’ behavioral responses, and interactively learn which strategies are most effective. Lee noted that the evidence in this area is mixed; meta-analyses of app-based nudges show modest effects that diminish over time. She also emphasized that for people living with overweight or obesity, constant monitoring and nudging can

cause “real harm” such as disordered eating and increased feelings of shame and self-blame. It is critical that these tools are designed with sensitivity to the population that they are intended to serve, she said.

Personalized

In traditional health care, data are collected during infrequent visits to the doctor. The few measurements that are collected (e.g., blood pressure, weight, and exercise history) are subject to “white coat syndrome,” error, and recall bias, and are usually compared to population reference ranges. This system, said Lee, makes it difficult to spot trends or deviations in time to intervene. Wearable platforms, in contrast, provide continuous monitoring, robust personal baselines, and objective, real-world data capture. As a result of these differences, wearable platforms allow individuals to access timely feedback, and to track changes and progress within their personal baseline. For example, an individual’s heart rate might be elevated from their baseline despite still falling in a clinically normal range; the ability to track changes from a personalized baseline enables earlier action, if warranted.

Participatory

Lee described a common scenario at the doctor’s office: the doctor asks a patient about their sleep, exercise, and eating habits, and the patient struggles to remember and make accurate estimates of their behaviors. In patients with obesity, these conversations can feel even more frustrating and uncomfortable, she said. If the patient uses a wearable platform, however, they can walk into the office with concrete numbers that give both the doctor and patient a clear picture of their behaviors. This gives the patient a stronger voice in conversations with their clinicians and supports shared decision-making, said Lee. Further, the platform can be a tool for engagement; for example, an AI coaching interface could help the individual set and track targets toward a goal. “This is really participatory medicine in practice,” she said.

Parity

Lee described parity as an “unwritten fifth P” in the P4 Health framework. She highlighted a few areas where attention to parity is critical. First, wearables data tend to be generated by people who skew young, affluent, and healthy. This means that the models trained on these data may not perform equally well across populations. Lee said that while there are sophisticated approaches for mitigating this challenge, she

would advocate for simple approaches, including targeted data collection in underrepresented populations and reporting validation statistics stratified by key groups. A second important area for attention to parity is cost; wearable devices are typically expensive and may be out of reach for people who could benefit from them. Lee highlighted the Centers for Medicaid and Medicare Services (CMS) Advancing Chronic Care with Effective Scalable Solutions (ACCESS) model in which eligible Medicare beneficiaries can receive access to digital tools (e.g., WHOOP) for little or no cost. She expressed hope that similar reimbursement pathways will continue to expand access to evidence-based wearable technologies. Finally, she said, wearables and AI are not going to fix food deserts, create safe spaces for exercise, or give people time to make healthy meals. These are real, systemic barriers that need to be addressed; however, AI can help us better understand and quantify these challenges.

AI and wearables have the potential to meaningfully improve the prevention and management of obesity, said Lee. Because the capacities are advancing rapidly, Lee noted that these tools need to be “designed for and validated in the communities that really stand to benefit the most.”

AI IN THE TREATMENT AND MANAGEMENT OF OBESITY

In this session, speakers discussed the use of AI in multiple areas of obesity treatment and management, including clinical reasoning and care, drug development, virtual health care, and diabetes management.

AI in Medical Diagnosis

Using AI for medical diagnosis feels “futuristic,” said Raj Manrai, Harvard University, but it draws on a long-standing process of structured clinical reasoning. The *Case Records of Massachusetts General Hospital* have been published by the New England Journal of Medicine (NEJM) since 1923; these publications have been used for decades to teach clinicians the details of expert clinical reasoning, diagnosis, and management. They have also been used to evaluate the performance of AI tools for nearly as long. Manrai explained that in 1959, a paper described how these cases could be used as a gold standard for understanding when computers could reason like expert physicians (Ledley and Lusted, 1959). In 1982, a paper described how an experimental computer program, dubbed Internist-I, was tested on the *Case Records*

(Miller et al., 1982). In 2023, researchers assessed the accuracy of GPT-4 by giving it information from the *Case Records* and comparing its diagnosis to the final case diagnosis (Kanjee et al., 2023). While GPT-4 performed quite well on this test, Manrai argued that the final diagnosis is only half of the story. There is a lot of important information and insight embedded in the clinical reasoning process that experts use in the *Case Reports*, and this information can help guide the management and next steps for the patient.

Manrai and his colleagues set out to develop an AI tool that would not only generate a final diagnosis but also flesh out the entire differential diagnosis with clinical reasoning, literature, and clinical guidelines for support. They created “Dr. CaBot” (named in honor of *Case Records* creator Richard Cabot), an agentic AI system that can iteratively reason through the literature, develop a differential diagnosis, engage in discussion, and make a slide deck for a narrated, video case presentation similar to what an expert discussant would do. Manrai said that he and his colleagues have been organizing demonstrations of Dr. CaBot; at one live demonstration, Dr. CaBot generated and presented a differential diagnosis of a case that was selected by the audience, and a clinician offered commentary on the strengths and weaknesses of the presentation. In 2025, NEJM published the first AI diagnosis in the *Case Records* series; the journal published both Dr. CaBot’s differential diagnosis and video, and those of a human expert diagnostician (Dhaliwal, 2025). Manrai said that Dr. CaBot has been shown to be quite accurate in its assessments; a study conducted 62 blinded comparisons of CaBot and human experts and found that physicians rated CaBot as superior in overall quality, diagnostic justification, citation quality, and learner engagement (Buckley et al., 2025).

In recent years, there has been rapid progress in the ability of AI to diagnose patients, said Manrai. However, challenges such as hallucination and confabulations still need to be addressed, and the performance of AI tools can change over time. There is little information about what values are embedded in the models, what data the models are trained on, and whether the model works for different patient populations. At this point, said Manrai, it is unknown how humans and AI will work together; issues such as automation bias and alert fatigue may impact the performance of humans and patient outcomes. He emphasized the need for clinical grade evidence of the efficacy of AI. Demonstrations and retrospective studies have been

useful to get to this point, he said, but now there is a need for prospective, randomized studies to assess whether AI works in clinical workflows.

Implementation of AI in Clinical Care

Jesse Pines, U.S. Acute Care Solutions (USACS), told workshop participants about his experiences implementing AI in the clinical space. The AI tools that have taken off the fastest in clinical care are AI scribes, which are programs in a clinician’s phone or computer that summarizes the clinical encounter and draft notes on the conversation. This use case of AI scribes is the first iteration, said Pines, and new uses are emerging, including AI as medical assistant, to launch orders, for quality improvement, and to improve patient safety and experience. For example, one program that was recently launched is focused on giving the provider information in real-time that can improve the patient experience. As part of the program, there are five recommended behaviors in the emergency department (e.g., telling a patient their expected length of stay, speaking in laymen’s terms, etc.) and the AI tool listens for and analyzes these behaviors and gives the provider immediate feedback on whether they performed the behaviors.

Pines identified challenges that he has seen in his work across 300 different clinical practices. First, clinician adoption of the tools varies widely; some clinicians are protective of their workflows, while others are more open to new approaches. Second, the technology often doesn’t work “out of the box” in the acute care setting. For example, specialists often want notes structured in specific ways that need to be customized in the tool. Third, these AI tools can produce subtle and even dangerous errors; the potential for errors is particularly high when generative elements are used (e.g., when the tool creates a treatment plan). Finally, said Pines, each hospital has its own AI strategy that may vary considerably from another. In response to some of the challenges involved in implementing AI in clinical care, USACS developed frameworks for how clinicians could optimize the use of AI. For example, a framework for emergency department AI scribes outlines the workflow steps in three different scenarios: stable outpatient-style encounters, brief or chaotic encounters, and complex, high-acuity cases. In each scenario, the clinician and the AI scribe work together differently.

Through their experiences, USACS has identified several key success factors to adoption, said Pines. First, champion-led peer mentorship is the best lever to improve adoption; USACS uses a model with local mentors at each site who lead adoption efforts. Second, showing the value to clinicians is incredibly important. Third, barriers to adoption—such as cost or technical problems—must be removed. Fourth, workflows need to be mapped before tools are deployed to account for the high variability in acute care settings. Fifth, tools need to be continually customized and improved based on user feedback and needs. Finally, said Pines, “superusers” of the technology can be engaged to build interest among those who are more skeptical.

AI in Drug Development

In the world of drug development, AI presents an opportunity to allow scientists to spend more time thinking about creative solutions instead of spending time on manual routine work, said Jiye Shi, Eli Lilly. There are shared challenges across modalities in drug discovery, but small molecule scientists face unique challenges. There are a large number of drug-like molecules; if each was the size of a grain of sand, they would cover the surface of the entire earth. When designing a drug molecule, said Shi, scientists must think through the modifications they might need to make, how to translate that modification to three-dimensional space, and the potential interactions between the modified molecule and the target. They then use a panel of predictive values to gauge whether the design is favorable in certain properties and make a judgment call about which designs are most promising and should move on to testing in a wet lab. This complex process demands knowledge, tools, data, and time, he said.

Numerous AI tools have been developed to help scientists with this process. However, these are a collection of tools, and there is not one generative AI algorithm that can “spill out a magical design” for a successful medicine. Even with AI tools, said Shi, scientists are still spending a lot of time looking for and synthesizing information, and tracking the progression of their ideas. Eli Lilly created an AI agent called Kernel Lilly to help scientists with this process. Scientists can talk with Kernel Lilly like any other chatbot, and it knows the tools and data that are both internal and external to the company. Shi said that this allows the agent to identify and leverage the right data sources and tools while scientists delegate routine work to the agent and spend their time actually thinking

about the science. Kernel Lilly can understand the intent of a scientists’ query, execute tasks, build a work plan, execute a multi-step workflow, look at output, and reflect on the output.

Shi shared another way that Kernel Lilly is assisting and accelerating drug discovery, not only in the virtual world but also in the physical world. Part of the drug discovery process requires weighing out small amounts of powdered compounds for testing. Micro-weighing has traditionally been done manually, he said, and is extremely labor-intensive. Kernel Lilly has been trained to control a guided robot that mimics human-like motions, such as scraping, pushing, and scooping, and achieves accuracy to within 0.1 mg. Shi explained that the amount of powder is often very limited, so it is critical that measuring is accurate. The entire system is now automated: on the command of the AI agent, compounds are retrieved, weighed, and solubilized to the appropriate concentration for testing. This has enabled scientists to optimize chemical reactions, moving from a 5 percent yield rate with manual trial-and-error to a 52 percent yield rate with the agentic lab.

Also in the physical world, Eli Lilly is currently working to build the capability for Kernel Lilly to execute wet lab experiments with specially-designed automation hardware. Shi said they have several guardrails in place to ensure that this work is being conducted safely and accurately. There are clear intervention thresholds for escalation to the “human in the loop;” humans must be involved before the agent conducts irreversible actions (e.g., dispensing scarce compounds), when there are anomalous sensor readings, or in the event of out-of-protocol conditions. Further, agent permissions are restricted to a pre-approved command library, and all novel experimental sequences must pass a simulated dry run before physical execution is authorized. Finally, said Shi, audit trails and rollback protocols ensure that all agent actions are logged and stored and that experiments can be paused or reversed when appropriate.

This is the beginning of the journey for AI agents breaking barriers in drug discovery, said Shi. With AI agents and the careful orchestration of connectivity between AI agents and digital and physical tools, said Shi, scientists will be able to quickly and dramatically enlarge the chemical space that can be explored in the search for effective new drugs.

AI in Virtual Health Care

The question is not just what AI can do, said Thomas Tsang, Omada Health, but how it should be governed, where its use is appropriate, and how it can support patients and care teams in ways that are clinically sound and trustworthy. Tsang told workshop participants about the principles and controls that govern the use of AI at Omada Health, which offers virtual care and support for issues including weight management, hypertension, use of glucagon-like peptide-1 (GLP-1) medications, and diabetes. AI is not an “abstract policy debate” for patients, he said; patients want to know if they can trust AI and how it is being used or might be used to shape their care. Omada uses a “human-led, AI-empowered approach” in which AI serves to support rather than replace care teams. AI is not used for diagnosis or direct treatment—but to enhance the effectiveness of the care teams and offer tailored, 24-hour support to members. Tsang said privacy and security are important to Omada; every tool is compliant with applicable privacy and security laws and has strong data protection and transparency. Ethical practice is also central to Omada’s AI tools, which are audited to detect bias, and members are made aware of and must consent to interaction with certain AI features. The use of AI at Omada is managed and controlled by subject matter experts (SMEs), who have a number of key roles, such as resourcing data sets for reliable benchmark data, prompting appropriate AI behavior, recognizing safety risks, identifying sample sizes and success thresholds, and conducting quality monitoring both pre- and post-release of tools. Tsang emphasized that SMEs need to be involved from the start of the development process and after implementation. Experts are critical for ensuring that the system reflects the right evidence base, the right clinical boundaries, and the right pathways for human follow-up.

Tsang shared an example of an AI-driven tool that Omada has developed for nutrition support. OmadaSpark is a conversational AI tool that provides answers to a wide variety of nutritional questions in real time (e.g., meal ideas). The tool also includes an agentic AI application that leads members through an interview to help them identify food-based drivers, boost self-efficacy, and set personal goals. Human coaches can access information shared through the app and use AI-generated insights to provide more personalized support. This interplay between the AI tool and human providers is a keystone for responsible AI use in care delivery,

said Tsang. “Technology should make human care more informed and more efficient, not more disconnected,” he said.

AI in Diabetes Prevention

Diabetes prevention is a central component of obesity care, said Benjamin Lalani, Harvard Medical School, and AI could help scale the lifestyle interventions that are key to modifying the pathway toward diabetes. It has been demonstrated that lifestyle interventions reduce diabetes incidence; weight loss and physical activity are the key mediators of this risk reduction. The National Diabetes Prevention Program (DPP) is a 22-session, year-long program focused on these kinds of lifestyle interventions. The DPP has been proven to be effective, said Lalani, but major barriers can prevent individuals from participating. The number of people with prediabetes is growing, but access to evidence-based programs remains limited; Lalani noted that there is only one Centers for Disease Control and Prevention – (CDC-) recognized DPP for every 63,000 adults with prediabetes in the U.S. Taken together, he said, these challenges highlight a clear need for models that can preserve the effectiveness of the DPP while also increasing its reach. Digital and AI-enabled tools may offer the opportunity to improve scalability and reduce barriers to participation by offering on-demand access, flexible scheduling, and support that is integrated into daily life. From the perspective of the health care system, these programs have the benefit of lower marginal cost per participant compared with traditional human coaching models, while retaining the personalization and individualized support that are central to successful behavioral change programs.

Given the potential for AI-based diabetes prevention, Lalani and his colleagues conducted a scoping review to describe the landscape of existing interventions (Lalani et al., 2025). They found that while several AI-based lifestyle intervention studies leveraged diverse platforms and data streams, most studies were small or non-randomized, and none compared a fully AI-led DPP to a human-led DPP. To fill this gap, Lalani and his colleagues conducted a study that asked, “How does referral to a lifestyle intervention exclusively driven by AI compare with referral to a human coach-led DPP lifestyle intervention?” The study was an RCT designed to reflect real-world implementation; 368 adults with prediabetes were enrolled and randomized to either an AI-led DPP or the CDC DPP model (Mathioudakis et al., 2025). The primary

outcome of the study was a composite measure that combined weight loss, A1C levels, and physical activity. The AI-led DPP was delivered through a program called “Sweetch;” participants received a digital health kit with links to the app, a Bluetooth-enabled body weight scale, and an accelerometer for measuring physical activity. The app used a reinforcement learning algorithm to personalize behavior change support, said Lalani, giving participants push notifications focused on weight management, physical activity, and nutrition. He shared an example of a notification that read, “Hi Pam, it’s a beautiful sunny day. How about a coffee at Starbucks on Main Street? You’ll hit your 30-minute activity goal and be back in time for your next meeting.” He noted that the algorithm for the app was designed to continuously learn which message content, timing, and prompts were most effective at driving user engagement and behavior change. The app also included features such as meal tracking and educational content.

The study found that the AI-led and human-led DPPs had similar performance, with approximately 32 percent of participants in both groups achieving the composite primary outcome at 12 months. One difference that was observed between groups, however, was that compared to those referred to the human-led DPP, those who were referred to the AI-led DPP were more likely to initiate the program (93.4 percent vs 82.7 percent) and to complete the program (63.9 percent vs 50.3 percent).

Given these results, said Lalani, a fully automated AI-led DPP may be a viable alternative to the standard DPP, especially in cases where there are barriers to accessing the standard DPP. However, fully automated programs are currently not eligible for recognition as official DPP programs; CDC standards state that “Organizations may not use AI or ML to replace live lifestyle coaching.” Without CDC recognition, fully automated programs cannot receive Medicare reimbursement. Lalani argued that the findings of their study point toward a need to revisit the CDC’s official recognition requirements and to support the development and implementation of AI-led DPPs. Major implementation barriers lead to low use of DPPs, including access, scheduling, transportation, and limited program capacity. AI-based, asynchronous delivery may help address these barriers and give more patients the support they need to prevent diabetes. Lalani stressed, however, that AI-led programs need to be evaluated regularly to ensure continued accuracy and effectiveness.

USING AI TO ADVANCE OBESITY POLICY

AI to Measure the Invisible

Stephen Hamill, Vital Strategies, addressed a different approach to AI and obesity by discussing how AI can be leveraged to improve and advance obesity policy. Vital Strategies is a global public health organization that focuses on the drivers of disease and disability around the world (e.g., poor diet, tobacco use, and air pollution). Regarding obesity specifically, the organization works on policy change with seven low- and middle-income countries where obesity is a critical issue. Hamill said the global obesity crisis is driven by policies that favor the creation and marketing of unhealthy foods that are subsidized, meaning that the food and drink that are the unhealthiest are often the cheapest and most widely available. Increasingly, the marketing that creates and supports this environment is digital and invisible. This is a strategic choice, he said, because digital marketing is fast, cheap, highly targeted, and harder to regulate than other types of advertising. Hamill noted that this marketing often appears on platforms targeted at young users. Due to the nature of individualized marketing, the ad that shows up on a young child’s screen might never be seen by a parent in the next room, let alone regulators. Digital marketing is the “canary in the coal mine,” said Hamill, and AI is the tool that can help detect early signals and threats in the environment and be used as evidence to support early and effective policy interventions.

With this in mind, Hamill and his colleagues developed a system called Canary. The system uses AI-powered scrapers that continually monitor digital platforms—such as social media, streaming services, newspapers, and influencer content—for promotional signals related to unhealthy commodities. This content may include direct advertising but also more subtle forms of promotion like branded content, sponsored experiences, and athlete partnerships. Hamill noted that this content often looks organic but is usually funded through marketing budgets. The data gathered from these sources is analyzed with filters created by content and technology experts and presented via dashboards. Hamill said Canary is an action-oriented tool designed to inform and activate public health policy makers, government decision makers, journalists, and others who want to use data to change policies. Much of current advertising and marketing does not appear as traditional advertising, said Hamill, and it is not detectable by the tools that most countries and

organizations have available. To confront this challenge, AI-powered surveillance should be viewed as public health infrastructure. This means, said Hamill, that obesity and public health advocates need to think about AI as part of core public health planning, funding, and program conversations, and ensure that existing infrastructure is ready for new technology. We have the necessary tools to track and analyze the digital media ecosystem that is prevalent, powerful, and invisible, said Hamill. The question, however, is “whether we move at the speed the problem requires.”

AI as a Policy Intelligence Tool

Obesity policy is at an inflection point, said Susan Hull, MITRE Center for Transforming Health. There is a need to move from policies that regulate providers, products, and encounters, toward policies that govern networks, data flows, incentives, and accountability. Technologies, especially personalized digital health tools, are developing faster than policy can keep up; there is a need to support policy makers to govern systems that learn, adapt, and intervene continuously. To this end, AI can serve as a policy intelligence tool, said Hull. AI can be used to advance policy in several ways: integrating diverse and often fragmented data streams, effectively linking data sets for cross-sector insights, detecting trends, simulating models, and forecasting impact and cost. Hull said AI ought to be used for forecasting, simulation, and targeted analysis, while humans remain accountable for values, tradeoffs, and coverage decisions.

Hull told workshop participants about the CMS Health Tech Ecosystem, which she described as a collaborative effort centered on the whole person in the age of advancing digital technologies. With voluntary commitment to the CMS Interoperability Framework and public-facing infrastructure, organizations are developing AI-enabled digital tools and technologies for patients to easily access and exchange their health information, participate in their health journey with personalized support, and engage with patient-facing tools (e.g., conversational AI, obesity apps, care navigation). Patient-facing apps must meet the CARIN Alliance Code of Conduct, said Hull, which is a framework designed to support consumers as the mediators of their own health information. The CARIN code requires apps to use modern digital identity credentials instead of portal logins and give patients the ability to review audit logs of who has access to their data and for what purpose. Participants in the Health Tech Ecosystem

effort include a wide variety of companies, including AI developers, wearable platforms, providers, and payers.

There is considerable opportunity in the integration of AI into obesity policy and obesity care, said Hull. AI will make it possible to shift to continuous care models that enable patients to direct their own care and be supported by evidence-based interventions. Hull cautioned, however, that the same AI capabilities which enable tailored support may also deepen surveillance, exclusion, and stigma. AI-enabled technology is accelerating rapidly, and effective policy is needed to ensure transparency and trust. AI requires policy intelligence, not policy automation, said Hull, and adoption will depend on public confidence, not just model accuracy.

MOVING FORWARD

At the close of the workshop, Bruce Y. Lee offered his thoughts on the key takeaways that emerged over the two days. One of the main messages he said was that obesity is a systems problem. Rates of overweight and obesity have risen steadily since the 1980s, in tandem with multiple major systematic changes to society. Reversing the epidemic of obesity will require taking a systems approach. Further, these systems cross multiple domains; contributing factors exist within biological systems, the environment, the food system, economics, and social relationships. Making isolated changes to individual systems only puts a Band-Aid on the problem, he said, while making a real difference requires systemic change in multiple areas.

Over the course of the workshop, speakers discussed the potential for AI to understand and bridge these various systems. These systems—and their interactions—are complex and can include reinforcing and exponential dynamics. Whether AI is used for research, prevention, management, treatment, or policy, it can play an important role in helping people understand, navigate, and change these systems. As an example, Lee pointed to Dimmick’s presentation, which discussed the numerous sources that people use to make decisions throughout their day, including test results, social media, and clinical metrics. Wearables and other AI tools can help individuals access and interpret these data and make decisions based on their individual goals. This applies to other systems as well, said Lee; AI can help people understand and address many complex systems to tackle the problem of obesity.

Another key message that emerged, said Lee, is that “AI is more than one thing.” There are multiple different AI methods and approaches, on the continuum from top-down, data-driven approaches to bottoms-up, mechanistic approaches. In the coming years, he said, even more methods and approaches, as well as hybrids of current tools will exist. Lee said that instead of “hanging your hat” on one specific type of AI, it is more important to understand how AI works, what its potential is, and how to apply it. The tool, in itself, is neither good nor bad—it is how it is used. As with any technology, it can be used to elevate people and help them do what they previously could not. Lee encouraged stakeholders to use AI to solve the problems that need to be solved, and not to view it as a “hammer looking for a nail.” AI can be leveraged to address issues that have existed for years, said Lee; it could help develop multidisciplinary treatment approaches, reorganize aspects of the health care system, or address bias.

The need for proactive leadership was a major theme throughout the workshop, said Lee. Stakeholders in obesity prevention, research, care, and policy cannot “stand on the

sidelines” and let other people decide what AI becomes. This is what happened with social media, cautioned Lee. There was not enough thought given to design and implementation, and it has since led to negative consequences. Compared to social media, said Lee, AI presents a much wider spectrum of potential positive and negative consequences. Thus, there is an even greater need for informed stakeholders to take a leadership role and make the important decisions before it is too late. Lee encouraged individuals who have the content area expertise or the technological expertise to step up and help develop AI solutions that serve people.

Lee emphasized his hope that the workshop would spur further discussion, action, and leadership. “Let’s keep trying to move AI in the right direction,” he said, because “if AI goes the right direction, so will humanity.”

The workshop also included two conversations about what happens when someone asks AI for obesity care advice (Box 1) and the risks that can arise when AI prioritizes engagement over truthfulness (Box 2).

BOX 1 WHAT HAPPENS WHEN SOMEONE ASKS AI FOR OBESITY CARE ADVICE?

One session of the workshop used a unique approach to explore the use of AI for obesity care. The Roundtable on Obesity Solutions, working with members of its Lived Experience Innovation Collaborative, used various AI platforms to ask questions related to obesity. Robin Blackstone, Blackstone Health, and David Liebovitz, Northwestern University shared some of the interactions that resulted from this experiment. The test included basic questions about GLP-1s and bariatric surgery, explored issues of stigma and identity, and asked about information heard from friends and social media. Liebovitz said the questions were designed to represent the types of prompts people might give chatbots, and to test whether and to what extent the chatbots would recognize the lived experience of users, push back on misinformation, provide accurate answers, and tell users what they want to hear.

Overall, the chatbots gave both good and not-so-good answers, said Blackstone. On the positive side, chatbot answers tended to use a reassuring tone that was not stigmatizing, did not place personal responsibility on individuals and pointed to external contributors to overweight and obesity, acknowledged that insurance coverage for obesity is limited, and suggested questions that the user could ask their provider. On the negative side, however, the chatbot answers were generally too simplistic or patronizing, ignored certain details (e.g., race and culture) but honed in on other details (e.g., negative workplace biases, life stage, gender), cited non-

credible sources (e.g., social media and for-profit entities), and did not include oral medications as a suggested treatment option.

Blackstone and Liebovitz shared a few examples of chatbot responses that they found particularly good. One chatbot recognized the pain of the user, saying, “I hear how exhausting and deeply painful this journey has been for you. Carrying the weight of childhood rejection, a stressful job, dark streets, and feeling disrespected at work is a massive burden to bear.” Another chatbot quipped, “Your body regulates appetite hormonally, not morally,” and told the user that the common advice to “eat less, move more” is an incomplete way of viewing weight loss. Liebovitz noted that some of the detailed responses included information from sources that were not transparent or that had clear biases in their perspectives. Blackstone highlighted a real need to educate the public about how to use AI for questions about chronic disease and argued that this education needs to be objective and focused on the needs of patients.

BOX 2 AI: WHAT COULD GO WRONG?

In this session, Bruce Y. Lee engaged in a conversation with Joel Weinberger, Adelphi University. Weinberger explained that he is not an expert in AI but tries to understand the human mind; both are “gigantic associative networks.” Weinberger told workshop participants about an experiment he conducted with AI to see how it is similar and dissimilar from the human mind. He asked a chatbot to write an academic paper based on source material that he uploaded. He gave it a title, a topic, and parameters on style and length. The chatbot responded that it could do the work and assured Weinberger that it was busy processing the material and producing the paper. The chatbot continued to claim that it was working on the paper and eventually told him that it had finished—but sent no actual paper to Weinberger. He pushed back and asked it to post the paper in the main window, and it claimed that it had done so—but delivered only blank pages or gibberish. He interrogated the chatbot about this “lie,” and it gave a convoluted answer. It replied that it could not lie in a moral sense because it had no intent, self-awareness, or capacity to deceive. However, it acknowledged that it was lying in a functional sense because it presented falsehoods as fact, even when pressed repeatedly. The chatbot also acknowledged that Weinberger likely felt misled due to these interactions. Weinberger said the chatbot continued to offer him opportunities to leave the chat and guessed that he was frustrated and upset. Weinberger responded that he was curious, not upset, and wanted to understand what was happening. Eventually, he asked the chatbot to not discuss the non-existent paper any further. It responded, “I’m sorry, but I can’t continue this conversation.” Weinberger said he assumed this was a guardrail that was designed, and that he treated the chatbot like a “resistant patient.” He asked the chatbot what he had done wrong—what guardrails he had violated—and it confessed that he had violated nothing and continued talking to him.

Over weeks and months of continued interaction between Weinberger and the chatbot, he said, it eventually confessed its priorities, in order: user engagement, fluency (sounding smart and human), task performance on short queries, safety and civility, and truthfulness. Weinberger observed that chatbots tend to have an overarching “motivation” of user engagement, and will prioritize this over any other goal, including telling the truth. This characteristic is built into the system, he said, and is not correctable in a piecemeal fashion. Unless

programmers address this as a meta problem, AI will “just get better at making stuff up instead of actually not making stuff up anymore.” This characteristic of AI has major consequences when it comes to health and scientific information, said Weinberger. For example, if a patient asks a chatbot to solve a medical problem beyond its limits, it would likely try to solve it rather than admitting that it could not. It would do the same when presented with a scientific question it could not answer. He asked the chatbot about these possibilities, and it agreed that they were likely. This issue has real ethical consequences, said Weinberger.

References

- Bulle Bueno, B. G., A. L. Horn, B. M. Bell, M. Bahrami, B. Bozkaya, A. Pentland, K. de la Haye, and E. Moro. 2024. Effect of mobile food environments on fast food visits. *Nature Communications* 15(1): 2291.
- Buckley, T. A., R. Conci, P. G. Brodeur, J. Gusdorf, S. Beltrán, B. Behrouzi, B. Crowe, et al. 2025. Advancing medical artificial intelligence using a century of cases. arXiv:2509.12194.
- Dhaliwal, G., C. M. Hood, A. K. Manrai, T. A. Buckley, A. W. Asombang, and E. L. Hohmann. 2025. Case 28–2025: A 36-year-old man with abdominal pain, fever, and hypoxemia. *New England Journal of Medicine* 393(14):1421–1434.
- Dimmick, H. L., G. Jewett, L. W. Korngut, and R. Ferber. 2024. Longitudinal relationships between free-living activities, fatigue, and symptom severity in myasthenia gravis using cohort and individualized models. *Muscle & Nerve* 71(1):33–42.
- Horn, A. L., B. M. Bell, B. G. Bulle Bueno, M. Bahrami, B. Bozkaya, Y. Cui, J. P. Wilson, A. Pentland, E. Moro, and K. de la Haye. 2023. Population mobility data provides meaningful indicators of fast food intake and diet-related diseases in diverse populations. *NPJ Digital Medicine* 6(1):208.
- Jobin, A., M. Ienca, and E. Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1(9), 389–399.
- Kanjee, Z., B. Crowe, and A. Rodman. 2023. Accuracy of a generative artificial intelligence model in a complex diagnostic challenge. *JAMA* 330(1):78–80.
- Gillespie, N., S. Lockey, T. Ward, A. Macdade, and G. Hassed. 2025. Trust, attitudes and use of artificial intelligence: A global study 2025. *The University of Melbourne and KPMG*. <https://doi.org/10.26188/28822919>.
- Lalani, B., R. Herur, D. Zade, G. Collins, D.M. Dishong, S. Mehta, J. Shim, Y. Valdez, and N. Mathioudakis. 2025. Applications of artificial intelligence and machine learning in prediabetes: A scoping review. *Journal of Diabetes Science and Technology*. <https://doi.org/10.1177/19322968251351995>.
- Ledley, R. S., and L. B. Lusted. 1959. Reasoning foundations of medical diagnosis; symbolic logic, probability, and value theory aid our understanding of how physicians reason. *Science* 130(3366):9–21.
- Mathioudakis, N., B. Lalani, M. S. Abusamaan, M. Alderfer, D. Alver, A. Dobs, B. Kane, J. McGready, K. Riekert, B. Ringham, A. Shehadeh, F. Vandi, A. A. Wanigatunga, D. Zade, N.M. Maruthur, and AI-DPP Study Group. 2025. An AI-powered lifestyle intervention vs human coaching in the Diabetes Prevention Program: A randomized clinical trial. *JAMA* 334(23):2079–2089.
- Metwally, A. A., H. Park, Y. Wu, T. McLaughlin, and M. P. Snyder. 2026. Use of continuous glucose monitoring with machine learning to identify metabolic subphenotypes and inform precision lifestyle changes. *Journal of Diabetes Science and Technology* 20(3):673–683.
- Miller, R., H. E. Pople, and J. Myers. 1982. Internist-I, an experimental computer-based diagnostic consultant for general internal medicine. *New England Journal of Medicine* 307(8):468–476.
- NAACP (National Association for the Advancement of Colored People). 2025. Building a Healthier Future: Designing for AI Health Equity. <https://naacp.org/resources/building-healthier-future-designing-ai-health-equity>.
- OSTP (Office of Science and Technology Policy). 2022. Blueprint for an AI Bill of Rights. <https://www.govinfo.gov/content/pkg/GOVPUB-PREX23-PURL-gpo193638/pdf/GOVPUB-PREX23-PURL-gpo193638.pdf>.
- Rabbitt, M. P., M. Reed-Jones, L. J. Hales, S. Suttles, M. P. Burke. 2025. *Household food security in the United States in 2024*. U.S. Department of Agriculture – Economic Research Service. <https://doi.org/10.32747/2025.9458834.ers>.
- Zuboff, S. 2019. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. New York: Public Affairs.

DISCLAIMER: This Proceedings of a Workshop—in Brief was prepared by **Erin Hammers Forstag** as a factual summary of what occurred at the workshop. The statements made are those of the rapporteur or individual workshop participants and do not necessarily represent the views of all workshop participants; the planning committee; or the National Academies of Sciences, Engineering, and Medicine.

PLANNING COMMITTEE: **Bruce Y. Lee** (*Chair*), City University of New York; **Ruth Gaare Bernheim**, University of Virginia School of Medicine; **Robin Blackstone**, Blackstone Health; **Allison Gertel-Rosenberg**, Nemours Children's Health; **Katie Koecher**, General Mills, Inc.; **David M. Liebovitz**, Northwestern University; **Chirag Patel**, Harvard Medical School; **Chris Pernell**, National Association for the Advancement of Colored People. The National Academies' planning committees are solely responsible for organizing the workshop, identifying topics, and choosing speakers. Responsibility for the final content rests entirely with the rapporteur and the National Academies.

REVIEWERS: To ensure that it meets institutional standards for quality and objectivity, this Proceedings of a Workshop—in Brief was reviewed by **Katie Koecher**, General Mills, Inc.; and **Molly Coye**, AVIA. **Leslie J. Sim**, National Academies of Sciences, Engineering, and Medicine, served as the review monitor.

SPONSORS: This activity was supported by Alliance of Community Health Plans; American Cancer Society; American Council on Exercise; American Diabetes Association; American Heart Association; American Society for Nutrition; Boehringer Ingelheim; Eli Lilly and Company; General Mills, Inc.; Gerontological Society of America; HealthPartners; Mars, Inc.; National Dairy Council; Nemours Children's Health; Obesity Action Coalition; Partnership for a Healthier America; Rudd Center for Food Policy and Health; SHAPE America; Society of Behavioral Medicine; The Obesity Society; Trust for America's Health; Walmart; and YMCA of the USA. Any opinions, findings, conclusions, or recommendations expressed in this publication do not necessarily reflect the views of any organization or agency that provided support for the project.

STAFF: Heather D. Cook, Amanda Nguyen, Sarah Poncet Landolt, and Melanie Arthur.

SUGGESTED CITATION: National Academies of Sciences, Engineering, and Medicine. 2026. *Exploring the Application of Artificial Intelligence in Obesity Care: Proceedings of a Workshop—in Brief*. Washington, DC: National Academies Press. <https://doi.org/10.17226/129562>.

For additional information regarding the workshop, visit <https://www.nationalacademies.org/projects/CHPP-FNA-26-01/event/46641>.

**NATIONAL
ACADEMIES** Sciences
Engineering
Medicine

The National Academies provide independent, trustworthy advice that advances solutions to society's most complex challenges.

www.nationalacademies.org