*Advances in Consumer Research*
https://acr-journal.com/
Volume-3 | Issue-10 |Sep: 2026
*Original Researcher Article*

# From Predictive Analytics to Consumer-Centric Healthcare: Machine Learning and the Transformation of Health Services

**Amit Kumar[1]*, Rituraj Soni[2]**
[1]Assistant Professor, Department of CCSIT, Teerthanker Mahaveer University, Moradabad, Uttar Pradesh – 244001, India.
[2] Assistant Professor, Department of CSE, Engineering College, Bikaner, Rajasthan, India.
***Corresponding Author**
Amit Kumar
Assistant Professor, Department of CCSIT, Teerthanker Mahaveer University, Moradabad, Uttar Pradesh – 244001, India.
Email: amitpanwar889@gmail.com

**ABSTRACT**
Machine learning (ML) is shifting health services from retrospective, provider-centred analytics toward continuous, personalised and consumer-facing engagement, yet the consumer-research literature on this shift remains fragmented. This review integrates peer-reviewed clinical, information-systems and consumer-behaviour evidence with recent industry surveys to examine how ML is used along the consumer health journey and how consumers respond to it. The present scenario shows credible progress in image-based diagnosis, wearable-enabled screening, risk stratification, care navigation and conversational support, including large-scale evidence that a smartwatch algorithm can flag atrial fibrillation and that chatbot answers to patient questions can be rated higher in quality and empathy than physician replies in a public forum. It also shows persistent weaknesses: proxy-label bias that under-identifies Black patients for extra care, proprietary models that perform worse in external validation than reported, consumer resistance rooted in "uniqueness neglect", and low willingness to share health data with technology firms. The review proposes a conceptual framework in which a trust layer of consent, fairness, validation and human oversight mediates between prediction and consumer action. Future perspectives address generative AI, federated learning, adaptive interventions, consumer co-design and regulation, and a research agenda is set out for consumer researchers. The review concludes that consumer-centric ML will be judged less by predictive accuracy than by whether consumers understand, control and benefit from it.
**Keywords**: machine learning; consumer-centric healthcare; predictive analytics; algorithmic trust; digital health.

## 1. INTRODUCTION

Health services are being reorganised around the consumer. People who once received care through a single clinician now compare providers online, track their own physiology on personal devices, and expect digital interactions to match the quality of in-person care (Redpoint Global, 2023). In a 2024 McKinsey survey of US consumers, a quarter reported that they could not get the care they needed when they needed it (Cordina et al., 2024). Access and experience, not only clinical quality, now shape how health services are judged.

Machine learning (ML) is the technology most often proposed as a response. ML systems learn statistical patterns from large datasets and use them to predict, classify or generate outputs without being explicitly programmed for each task (Rajkomar et al., 2019). The first healthcare applications were predictive analytics built for institutions: identifying high-risk, high-cost patients for care management (Bates et al., 2014) or forecasting clinical events from electronic records (Obermeyer & Emanuel, 2016). The user was the health system, and the consumer was largely a data subject. Three developments have changed this. First, deep learning performed at or near specialist level on narrow perceptual tasks such as skin-lesion and retinal-image classification (Esteva et al., 2017; Gulshan et al., 2016). Second, consumer devices turned millions of people into continuous data sources (Perez et al., 2019). Third, large language models now let consumers interact with health information in natural language (Singhal et al., 2023; Thirunavukarasu et al., 2023). Together these shifts move ML from analytics that act on consumers toward systems that act with and for them (Topol, 2019).

We use consumer-centric healthcare to mean services designed around consumers' needs, preferences and capacity to participate, in which consumers have the information, choice and control to shape their own care. Consumer research has studied technology uptake through the technology acceptance model (Davis, 1989) and the extended unified theory of acceptance and use of technology, which adds hedonic motivation, price value and habit (Venkatesh et al., 2012). Reviews of wearable technology show how adoption depends on perceived usefulness, privacy concerns and sustained use (Ferreira et al., 2021). Health is, however, a high-stakes credence setting in which consumers cannot easily verify quality. Acceptance of ML there involves trust, fairness and identity as well as usefulness. Longoni et al. (2019)

showed across a series of studies that consumers are less willing to use healthcare delivered by AI than by human providers, because they fear that automated care neglects their unique circumstances.

Existing reviews tend to examine one side of the problem. Some focus on clinical performance (Esteva et al., 2019; Liu et al., 2019), others on ethics and implementation (Char et al., 2018; Vayena et al., 2018), and others on consumer opinion about diagnostic tools (Stephens et al., 2025) or on AI-based consumer health informatics (Asan et al., 2023). Few connect predictive capability to consumer experience, trust and behaviour in a single account.

This review therefore pursues four aims: (a) to map how ML is applied along the consumer health journey; (b) to synthesise evidence on consumer responses, including acceptance, trust, privacy and equity concerns; (c) to propose a framework that links prediction to consumer-centric outcomes; and (d) to outline future perspectives and a research agenda for consumer researchers. It is a narrative review. Sources were selected for relevance, influence and recency from the clinical, information-systems and consumer-research literatures. Industry surveys are included where peer-reviewed data are scarce and are identified as such. The article proceeds with the present scenario, future perspectives and a conclusion.

## 2. PRESENT SCENARIO

### 2.1 From prediction to consumer-centric action

The present scenario is best understood as a pipeline that has been extended beyond the clinic. Data from records, claims, wearables and consumer conversations feed ML models; the models produce predictions, classifications or generated text; and these outputs increasingly reach consumers directly as nudges, triage advice, navigation support or diagnostic information (Topol, 2019). Figure 1 summarises this pipeline and adds the element that the consumer-research evidence reviewed below shows to be decisive: a trust layer of consent and privacy, fairness, independent validation and human oversight. Clinical-impact research reaches the same conclusion from the supply side, arguing that the main barriers to value are not algorithms but evaluation, integration and trust (Kelly et al., 2019; Wiens et al., 2019).

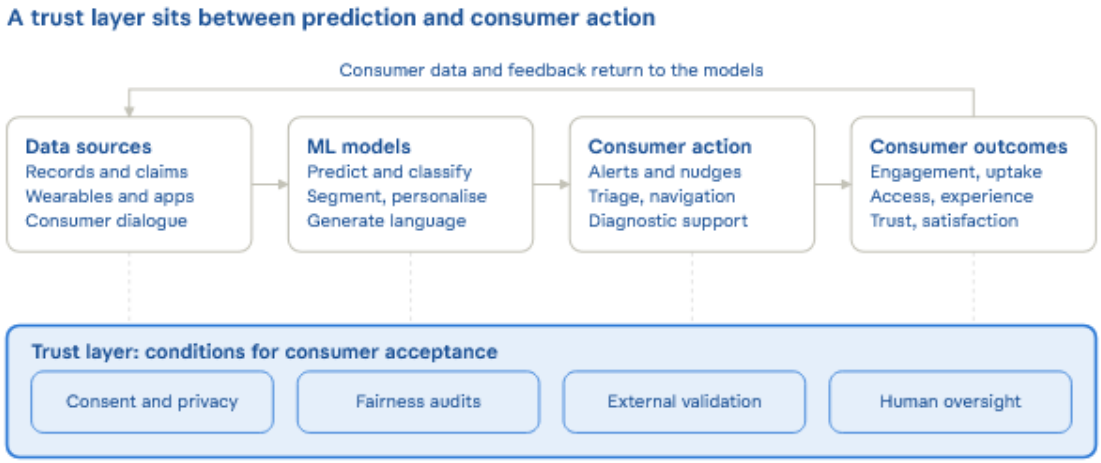



**Figure 1: Conceptual Framework Linking Machine Learning Predictions to Consumer-Centric Outcomes**

Note. Framework synthesised by the authors from the literature reviewed in this article. ML = machine learning.

### 2.2 Machine learning along the consumer health journey

Table 1 organises current applications by the stage of the consumer journey in which they operate. The evidence is uneven. Prevention, diagnosis and conversation have peer-reviewed support from large studies, whereas access and navigation rest mainly on industry reports.

**Table 1: Machine Learning Applications Along the Consumer Health Journey**

| Journey stage | Representative ML application | Illustrative evidence | Consumer-facing issue |
|---|---|---|---|
| Prevention and monitoring | Wearable-based screening and activity tracking | Irregular-pulse algorithm tested in 419,297 smartwatch users (Perez et al., 2019); trackers increase physical activity (Brickwood et al., 2019) | Sustained use, data accuracy, privacy (Ferreira et al., 2021) |
| Risk stratification | Predictive models for high-risk patients and clinical deterioration | Cost-based algorithm under-identified Black patients (Obermeyer et al., 2019); sepsis model missed 67% of cases in external validation (Wong et al., 2021) | Fairness, external validity, transparency |
| Diagnosis | Deep learning on skin and retinal images | Dermatologist-level lesion classification (Esteva et al., 2017); autonomous retinopathy detection in primary care (Abràmoff et al., 2018); pooled | Resistance to AI providers, accountability (Longoni et al., 2019) |

| | | | |
|---|---|---|---|
| | | comparison with clinicians (Liu et al., 2019) | |
| Access and navigation | Scheduling, no-show prediction, cost and coverage explanation | Industry-reported gains in one health centre; leaders see most generative-AI potential here (Cordina et al., 2024) | Consent control, data readiness, bias in outreach |
| Conversation and support | Chatbots, language-model replies, digital therapeutics | Systematic review of conversational agents (Laranjo et al., 2018); trial of a CBT chatbot (Fitzpatrick et al., 2017); chatbot replies rated higher for quality and empathy (Ayers et al., 2023) | Accuracy, perceived empathy, disclosure of AI use |

Note. ML = machine learning; CBT = cognitive behavioural therapy. Evidence entries are illustrative rather than exhaustive.

**Prevention and monitoring.** Consumer devices have made population-scale screening possible. In the Apple Heart Study, 419,297 participants without known atrial fibrillation were monitored over a median of 117 days, and 2,161 (0.52%) received an irregular-pulse notification (Perez et al., 2019). Among notified participants who returned usable ECG patches, 34% showed atrial fibrillation, and 57% of those who completed a follow-up survey contacted a health provider outside the study. The result is notable for consumer research because it shows an algorithmic alert prompting consumers to seek care on their own initiative. A meta-analysis found that consumer wearable activity trackers increased physical activity participation (Brickwood et al., 2019). Yet a systematic review of 74 publications on wearables and consumers found that research remains fragmented across health and marketing and that benefits are weighed against privacy threats to users and to others around them (Ferreira et al., 2021).

**Risk stratification.** Predictive analytics for high-risk, high-cost patients was the founding use of big data in health care (Bates et al., 2014). It also provides the clearest cautionary evidence. Obermeyer et al. (2019) studied a commercial algorithm used to allocate extra care and found that, at a given risk score, Black patients were considerably sicker than White patients, because the algorithm predicted health-care cost rather than health need. Correcting the label would raise the share of Black patients flagged for extra help from 17.7% to 46.5%. Wong et al. (2021) externally validated a widely implemented proprietary sepsis model in 27,697 patients and found poor discrimination: the model generated alerts for 18% of all hospitalised patients yet failed to identify most sepsis cases. Both cases involve decisions about consumers that the consumers themselves could not see.

**Diagnosis.** Deep learning has reached specialist-level performance on well-defined image tasks. A convolutional network trained on 129,450 clinical images classified skin lesions at a level comparable to dermatologists (Esteva et al., 2017), and a retinal-image algorithm achieved high sensitivity and specificity for diabetic retinopathy (Gulshan et al., 2016). A prospective pivotal trial then tested an autonomous system for diabetic retinopathy in primary care offices (Abràmoff et al., 2018). A systematic review and meta-analysis found diagnostic accuracy comparable to health-care professionals, while cautioning that few studies reported externally validated results or compared models and clinicians on the same samples (Liu et al., 2019). Regulatory clearance of AI-based devices has grown rapidly, which brings such tools within reach of consumers (Benjamens et al., 2020).

**Access and navigation.** In a McKinsey survey of healthcare leaders, 62% identified consumer engagement and experience as the area where generative AI has the greatest potential, yet only 29% had begun implementing generative AI for any purpose (Cordina et al., 2024). Applications include no-show prediction, personalised reminders and plain-language explanations of cost and coverage. These are industry claims and should be read as such; the same report notes that preparing integrated data accounts for roughly 70% of the work in building AI solutions. Outreach models also raise equity questions, because a patient predicted to miss appointments may be unreliable for reasons linked to disadvantage (Rajkomar et al., 2018).

**Conversation and support.** A systematic review found that conversational agents in health care were promising but heterogeneous in design and evaluation (Laranjo et al., 2018). A randomised trial of a fully automated cognitive behavioural therapy chatbot reported reduced depressive symptoms in young adults (Fitzpatrick et al., 2017), and a later meta-analysis found chatbot interventions effective in changing some lifestyle behaviours (Singh et al., 2023). Large language models extend this from scripted dialogue to open-ended conversation (Singhal et al., 2023; Thirunavukarasu et al., 2023). Ayers et al. (2023) compared physician replies with chatbot replies to 195 patient questions posted on a public forum. Licensed professionals preferred the chatbot in 78.6% of evaluations, rating its responses higher on both quality and empathy. The authors noted, however, that a forum exchange differs from a consultation, and the chatbot's replies were considerably longer.

### 2.3 Consumer acceptance and resistance

Capability does not guarantee uptake. In a series of experiments, Longoni et al. (2019) found that consumers were less likely to use healthcare delivered by AI, offered lower reservation prices for it, and were less sensitive to differences in provider performance. The mechanism was uniqueness neglect, the concern that an AI provider cannot account for one's own characteristics and

circumstances. Resistance was weaker when care was framed as personalised and was eliminated when AI supported rather than replaced a human decision-maker. A commentary re-examining the same data argued that consumers did prefer AI when it was explicitly more accurate than the human provider (Pezzo & Beckstead, 2020). Taken together, the studies suggest that resistance is conditional on perceived performance, personalisation and the presence of a human in the loop.

Qualitative evidence points the same way. A systematic review of 23 qualitative and mixed-methods studies identified four themes: trust, fear and uncertainty; data privacy and governance; impact on healthcare delivery and access; and a desire to be engaged in design and implementation (Stephens et al., 2025). Consumers' views of ML diagnosis depended on how far they trusted the clinicians using it. A scoping review of AI-based consumer health informatics similarly concluded that consumer-facing tools must be understood through patient perceptions and outcomes, not technical performance alone (Asan et al., 2023). Industry survey data are consistent: in one commissioned US survey, 62% of consumers expected online communications to match the in-person experience (Redpoint Global, 2023). Classic acceptance models (Davis, 1989; Venkatesh et al., 2012) capture usefulness, ease and habit, but they say little about uniqueness, accountability or algorithmic fairness, which the health evidence shows to be central.

### 2.4 Privacy, equity and the quality of evidence

Trust in the data custodian shapes whether consumers will supply the data on which ML depends. In McKinsey's 2024 consumer survey, willingness to grant access to personal health information was highest for primary care providers and lowest for large technology firms (Figure 2). Privacy concerns appeared in nearly all of the studies synthesised by Stephens et al. (2025), and the McKinsey authors note that consumers have no easy way to review or adjust consents already given (Cordina et al., 2024). Scholars of medical big data argue that conventional consent and de-identification are poorly matched to the reuse of data for ML (Price & Cohen, 2019).

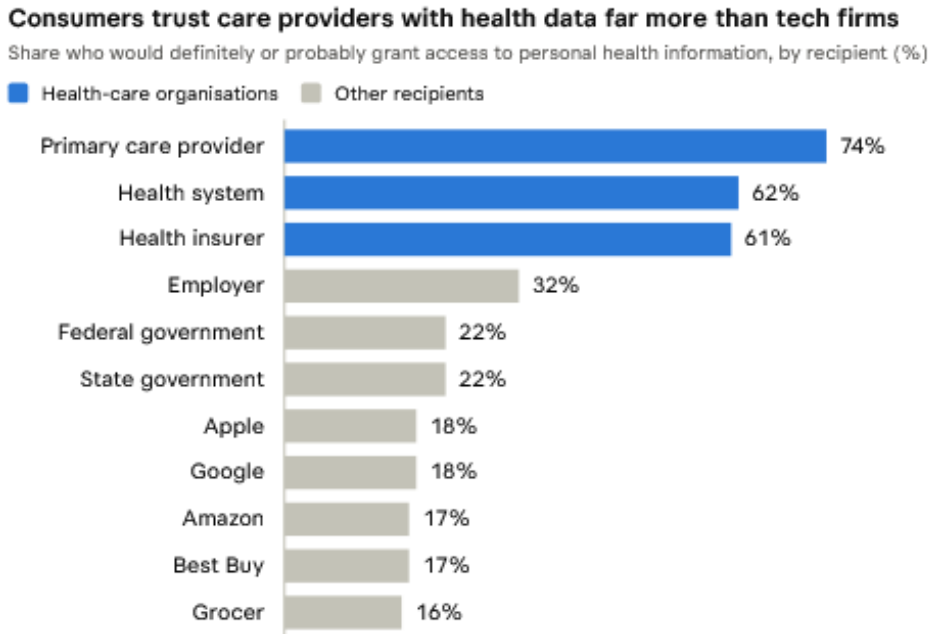



**Figure 2: Consumers' Willingness to Grant Access to Personal Health Information, by Recipient**

Note. Percentages of respondents who would definitely or probably grant access (n = 1,219 US consumers). Data from the McKinsey Consumer Health Insights Survey, 2024, as reported by Cordina et al. (2024).

Equity and evidence quality are related concerns. Beyond the proxy-label bias documented by Obermeyer et al. (2019), Rajkomar et al. (2018) describe how bias can enter through unrepresentative data, measurement and deployment, and propose that fairness be evaluated as part of model development. Independent validation remains uncommon, and proprietary models can underperform their published specifications (Wong et al., 2021). Ethical analyses call for transparency, accountability and attention to consent (Char et al., 2018; Vayena et al., 2018), and the World Health Organization (2021) has issued guidance that places human autonomy, transparency and equity at the centre of AI for health. The present scenario is therefore one of rapid technical progress, uneven evidence and unresolved consumer trust.

## 3. FUTURE PERSPECTIVE

The next phase of ML in health services will be shaped less by new model architectures than by how well five conditions are met: conversational interfaces that consumers can trust, data practices that consumers can control, engagement that adapts to the individual, evidence that is independently validated, and governance that assigns responsibility. Each is discussed below, followed by a research agenda.

### 3.1 Generative AI as the consumer interface

Large language models are likely to become the default front door to health information, because they can respond to free-text queries without task-specific training

(Thirunavukarasu et al., 2023) and encode substantial clinical knowledge (Singhal et al., 2023). Healthcare leaders already expect value here: generative AI can make previously unusable unstructured data, such as call transcripts, available for personalisation (Cordina et al., 2024). The same properties create risks that consumers will bear first. Fluent but inaccurate answers are difficult for non-experts to detect, and evaluation on examination questions is a weak proxy for clinical aptitude (Thirunavukarasu et al., 2023). The most defensible near-term model is therefore assistive, as suggested by Ayers et al. (2023), who proposed chatbot-drafted replies that clinicians review and edit. This design also fits the consumer evidence, because resistance to medical AI is reduced when AI supports a human decision-maker (Longoni et al., 2019). Whether consumers want to be told that a message was machine-drafted, and how disclosure affects trust, is an open question.

### 3.2 Privacy-preserving and consumer-controlled data

Consumers' reluctance to share data with technology firms (Figure 2) implies that future systems must learn from data without centralising it. Federated learning trains models across institutions while keeping records local and is being developed for health applications (Rieke et al., 2020). Technical safeguards alone will not resolve consumer concerns. Participants in qualitative studies asked for opt-in consent, the ability to see how data are used, and clear lines of responsibility (Stephens et al., 2025), and scholars have argued for consent and governance models suited to continual data reuse (Price & Cohen, 2019). McKinsey authors anticipate that consumers will expect the same control over health data that they already have in other sectors, such as choosing which purchases train a recommendation model (Cordina et al., 2024). Interfaces that make data use visible and revocable are likely to become a source of competitive advantage as well as compliance.

### 3.3 Adaptive, personalised engagement

The uniqueness-neglect finding suggests that personalisation is not merely a marketing feature but a precondition of acceptance (Longoni et al., 2019). Just-in-time adaptive interventions provide a design framework: they use real-time data from sensors and self-report to decide when and how to deliver support, with explicit decision points, tailoring variables and intervention options (Nahum-Shani et al., 2018). Combined with wearables, such interventions could address the sustained-use problem identified in the consumer wearable literature (Ferreira et al., 2021). Their promise depends on evidence that adaptation improves outcomes without increasing intrusiveness, and on transparent communication of why a particular prompt was sent.

### 3.4 Validation, fairness and regulation

Future deployments will need continuing, independent evaluation rather than one-off approval. The sepsis-model case shows how a widely used proprietary tool can perform far below its documented accuracy (Wong et al., 2021), and the cost-proxy case shows how an accurate prediction can still be unjust (Obermeyer et al., 2019). Practical roadmaps recommend that developers define the problem with stakeholders, assess data quality and bias, and monitor models after deployment (Wiens et al., 2019). Fairness assessment should be treated as part of model development (Rajkomar et al., 2018), and implementation pathways should address the translational gap between algorithmic accuracy and clinical impact (Kelly et al., 2019). Regulation is evolving alongside the growing catalogue of cleared AI devices (Benjamens et al., 2020), and international guidance now emphasises transparency, accountability and equity (World Health Organization, 2021). Consumers in qualitative studies also asked who is responsible when an algorithm errs, and found no clear answer (Stephens et al., 2025), which makes liability rules an unresolved consumer issue.

### 3.5 Implications and research agenda

For managers and policymakers, three principles follow from the evidence. ML services should be visibly personalised, so that consumers can see how their circumstances shaped a recommendation. They should keep a clinician or other accountable human in the loop where stakes are high. They should give consumers meaningful control over data use. For consumer researchers, the field offers a natural laboratory in which classic questions about trust, resistance, privacy and self-control are tested under high stakes. Table 2 outlines six priority themes.

**Table 2: Research Agenda for Consumer-Centric Machine Learning in Health Services**

| Theme | Open question | Suggested approach | Anchoring literature |
|---|---|---|---|
| Trust and acceptance | Do human oversight and personalised framing sustain acceptance when ML is embedded in real care rather than vignettes? | Field experiments and longitudinal consumer panels | Longoni et al. (2019); Pezzo & Beckstead (2020) |
| Data sharing and control | How do consent interfaces and the identity of the data recipient shape willingness to share? | Choice experiments; trials of dynamic consent | Cordina et al. (2024); Price & Cohen (2019) |
| Equity and fairness | How do consumers perceive and respond to disclosed algorithmic bias? | Experiments with fairness disclosures; audits with community panels | Obermeyer et al. (2019); Rajkomar et al. (2018) |

| | | | |
|---|---|---|---|
| Sustained engagement | Which adaptive designs sustain wearable and chatbot use beyond initial novelty? | Micro-randomised trials of adaptive interventions | Nahum-Shani et al. (2018); Ferreira et al. (2021) |
| Generative AI communication | When do consumers prefer, trust or reject AI-generated clinical messages, and does disclosure matter? | Blinded comparisons in patient portals with disclosure manipulations | Ayers et al. (2023); Singhal et al. (2023) |
| Value and outcomes | Do consumer-centric ML services improve experience, cost and health outcomes? | Pragmatic trials with consumer-reported measures | Kelly et al. (2019); Asan et al. (2023) |

Note. The agenda is the authors' synthesis of the literature reviewed in this article.

Consumer involvement in design is itself a research priority. Consumers want to be engaged in the development and implementation of ML tools, and doing so is associated with better quality, person-centred care and trust (Stephens et al., 2025). Studies that treat consumers as co-designers, rather than as subjects of acceptance surveys, would address a persistent gap in the literature.

## 4. CONCLUSION

Machine learning has moved from institutional predictive analytics to consumer-facing services that screen, advise, converse and navigate. The evidence reviewed here shows genuine capability: smartwatch algorithms that prompt consumers to seek care (Perez et al., 2019), specialist-level image classification (Esteva et al., 2017), and chatbot answers that evaluators rated as more empathetic than physician replies in a public forum (Ayers et al., 2023). It also shows that capability is not the binding constraint. Algorithms have misallocated care through poorly chosen proxies (Obermeyer et al., 2019), proprietary models have underperformed their documentation (Wong et al., 2021), and consumers resist automated care that seems to ignore their uniqueness (Longoni et al., 2019) or entrusts their data to firms they do not trust (Cordina et al., 2024).

The framework proposed in this article places a trust layer of consent, fairness, validation and human oversight between prediction and consumer action. Its implication is that consumer-centric healthcare cannot be achieved by deploying better models alone. It requires services that are visibly personalised, that keep accountable humans involved where stakes are high, and that give consumers real control over data. The limits of this review should be acknowledged. It is a narrative synthesis rather than a systematic search, it draws on industry surveys whose methods are less transparent than peer-reviewed studies, and the field is moving quickly, particularly in generative AI. Consumer researchers are well placed to supply the missing evidence on trust, acceptance and sustained engagement, and to ensure that the transformation of health services is judged by whether consumers understand, control and benefit from it

## REFERENCES

1. Abràmoff, M. D., Lavin, P. T., Birch, M., Shah, N., & Folk, J. C. (2018). Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. npj Digital Medicine, 1, Article 39.
2. https://doi.org/10.1038/s41746-018-0040-6
3. Asan, O., Choi, E., & Wang, X. (2023). Artificial intelligence–based consumer health informatics application: Scoping review. Journal of Medical Internet Research, 25, e47260.
4. https://doi.org/10.2196/47260
5. Ayers, J. W., Poliak, A., Dredze, M., Leas, E. C., Zhu, Z., Kelley, J. B., Faix, D. J., Goodman, A. M., Longhurst, C. A., Hogarth, M., & Smith, D. M. (2023). Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. JAMA Internal Medicine, 183(6), 589–596.
6. https://doi.org/10.1001/jamainternmed.2023.1838
7. Bates, D. W., Saria, S., Ohno-Machado, L., Shah, A., & Escobar, G. (2014). Big data in health care: Using analytics to identify and manage high-risk and high-cost patients. Health Affairs, 33(7), 1123–1131. https://doi.org/10.1377/hlthaff.2014.0041
8. Benjamens, S., Dhunnoo, P., & Meskó, B. (2020). The state of artificial intelligence-based FDA
9. -approved medical devices and algorithms: An online database. npj Digital Medicine, 3, Article 118. https://doi.org/10.1038/s41746-020-00324-0
10. Brickwood, K.-J., Watson, G., O'Brien, J., & Williams, A. D. (2019). Consumer-based wearable activity trackers increase physical activity participation: Systematic review and meta-analysis. JMIR mHealth and uHealth, 7(4), e11819.
11. https://doi.org/10.2196/11819
12. Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing machine learning in health care — Addressing ethical challenges. New England Journal of Medicine, 378(11), 981–983.
13. https://doi.org/10.1056/NEJMp1714229
14. Cordina, J., Coronado, E., Williams, P., & Greenberg, S. (2024, November 15). Harnessing AI to reshape consumer experiences in healthcare. McKinsey & Company. https://www.mckinsey.com/industries/healthcare/our-insights/harnessing-ai-to-reshape-consumer-experiences-in-healthcare
15. Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Quarterly, 13(3), 319–340.
16. https://doi.org/10.2307/249008
17. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural

18. networks. Nature, 542(7639), 115–118.
19. https://doi.org/10.1038/nature21056
20. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., & Dean, J. (2019). A guide to deep learning in healthcare. Nature Medicine, 25(1), 24–29. https://doi.org/10.1038/s41591-018-0316-z
21. Ferreira, J. J., Fernandes, C. I., Rammal, H. G., & Veiga, P. M. (2021). Wearable technology and consumer interaction: A systematic review and research agenda. Computers in Human Behavior, 118, Article 106710.
22. https://doi.org/10.1016/j.chb.2021.106710
23. Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial. JMIR Mental Health, 4(2), e19. https://doi.org/10.2196/mental.7785
24. Gulshan, V., Peng, L., Coram, M., Stumpe, M. C., Wu, D., Narayanaswamy, A., ... Webster, D. R. (2016). Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. JAMA, 316(22), 2402–2410. https://doi.org/10.1001/jama.2016.17216
25. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. BMC Medicine, 17, Article 195. https://doi.org/10.1186/s12916-019-1426-2
26. Laranjo, L., Dunn, A. G., Tong, H. L., Kocaballi, A. B., Chen, J., Bashir, R., ... Coiera, E. (2018). Conversational agents in healthcare: A systematic review. Journal of the American Medical Informatics Association, 25(9), 1248–1258. https://doi.org/10.1093/jamia/ocy072
27. Liu, X., Faes, L., Kale, A. U., Wagner, S. K., Fu, D. J., Bruynseels, A., ... Denniston, A. K. (2019). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: A systematic review and meta-analysis. The Lancet Digital Health, 1(6), e271–e297. https://doi.org/10.1016/S2589-7500(19)30123-2
28. Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. Journal of Consumer Research, 46(4), 629–650. https://doi.org/10.1093/jcr/ucz013
29. Nahum-Shani, I., Smith, S. N., Spring, B. J., Collins, L. M., Witkiewitz, K., Tewari, A., & Murphy, S. A. (2018). Just-in-time adaptive interventions (JITAIs) in mobile health: Key components and design principles for ongoing health behavior support. Annals of Behavioral Medicine, 52(6), 446–462. https://doi.org/10.1007/s12160-016-9830-8
30. Obermeyer, Z., & Emanuel, E. J. (2016). Predicting the future — Big data, machine learning, and clinical medicine. New England Journal of Medicine, 375(13), 1216–1219.
31. https://doi.org/10.1056/NEJMp1606181
32. Obermeyer, Z., Powers, B., Vogeli, C., &
33. Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453.
34. https://doi.org/10.1126/science.aax2342
35. Perez, M. V., Mahaffey, K. W., Hedlin, H., Rumsfeld, J. S., Garcia, A., Ferris, T., ... Turakhia, M. P. (2019). Large-scale assessment of a smartwatch to identify atrial fibrillation. New England Journal of Medicine, 381(20), 1909–1917.
36. https://doi.org/10.1056/NEJMoa1901183
37. Pezzo, M. V., & Beckstead, J. W. (2020). Patients prefer artificial intelligence to a human provider, provided the AI is better than the human: A commentary on Longoni, Bonezzi and Morewedge (2019). Judgment and Decision Making, 15(3), 443–445.
38. Price, W. N., II, & Cohen, I. G. (2019). Privacy in the age of medical big data. Nature Medicine, 25(1), 37–43. https://doi.org/10.1038/s41591-018-0272-7
39. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. New England Journal of Medicine, 380(14), 1347–1358.
40. https://doi.org/10.1056/NEJMra1814259
41. Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring fairness in machine learning to advance health equity. Annals of Internal Medicine, 169(12), 866–872.
42. https://doi.org/10.7326/M18-1990
43. Redpoint Global. (2023). Healthcare consumer perspectives on digital engagement and AI: Focus on healthcare providers and health plans in the U.S. https://www.redpointglobal.com/wp-content/uploads/2023/09/Dynata_2023_Healthcare_Consumer_Perspectives_on_Digital_Engagement__AI_042024.pdf
44. Rieke, N., Hancox, J., Li, W., Milletarì, F., Roth, H. R., Albarqouni, S., ... Cardoso, M. J. (2020). The future of digital health with federated learning. npj Digital Medicine, 3, Article 119.
45. https://doi.org/10.1038/s41746-020-00323-1
46. Singh, B., Olds, T., Brinsley, J., Dumuid, D., Virgara, R., Matricciani, L., ... Maher, C. (2023). Systematic review and meta-analysis of the effectiveness of chatbots on lifestyle behaviours. npj Digital Medicine, 6, Article 118. https://doi.org/10.1038/s41746-023-00856-1
47. Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., ... Natarajan, V. (2023). Large language models encode clinical knowledge. Nature, 620(7972), 172–180. https://doi.org/10.1038/s41586-023-06291-2
48. Stephens, J. H., Northcott, C., Poirier, B. F., & Lewis, T. (2025). Consumer opinion on the use of machine learning in healthcare settings: A qualitative systematic review. Digital Health, 11.
49. https://doi.org/10.1177/20552076241288631
50. Thirunavukarasu, A. J., Ting, D. S. J., Elangovan, K., Gutierrez, L., Tan, T. F., & Ting, D. S. W. (2023). Large language models in medicine. Nature Medicine, 29(8), 1930–1940. https://doi.org/10.1038/s41591-023-02448-8

51. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 25(1), 44–56. https://doi.org/10.1038/s41591-018-0300-7
52. Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. PLOS Medicine, 15(11), Article e1002689. https://doi.org/10.1371/journal.pmed.1002689
53. Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. MIS Quarterly, 36(1), 157–178. https://doi.org/10.2307/41410412
54. Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., ... Goldenberg, A. (2019). Do no harm: A roadmap for responsible machine learning for health care. Nature Medicine, 25(9), 1337–1340. https://doi.org/10.1038/s41591-019-0548-6
55. Wong, A., Otles, E., Donnelly, J. P., Krumm, A., McCullough, J., DeTroyer-Cooley, O., Pestrue, J., Phillips, M., Konye, J., Penoza, C., Ghous, M., & Singh, K. (2021). External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Internal Medicine, 181(8), 1065–1070. https://doi.org/10.1001/jamainternmed.2021.2626
56. World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. https://www.who.int/publications/i/item/9789240029200