RESEARCH PAPER

Ethical AI Deployment in Healthcare Marketing: A Management Perspective on Computer Science-Enabled Patient-Centric Strategies

Dr Resmi R[1], kruthika C G[2], Dr. Nafeza. E[3], Dr. Gayathri. A[4], Dr. Sajan Jerome[5], Dr Anilkumar E N[6]

[1]Professor, Electronics and Communication Engineering,
LBS Institute of Technology for Women ,Poojappura.
Email: resmi@lbsitw.ac.in, ORCID id : https://orcid.org/0000-0002-0849-5509
[2]Assistant professor, Department of Artificial intelligence and Machine learning
Nitte Meenakshi Institute of Technology ,NITTE(Deeemed to be University)
Bangalore ,Karnataka , India 560064
Email: kruthika.cg@nmit.ac.in
[3]Assistant professor, MBA, Velammal Engineering College.
ORCID id: 0009-0005-7355-982X
[4]Head & Assistant professor, BBA, Anna Adarsh College for women.
Email: gayathria@annaadarsh.edu.in
[5]Associate Professor, Mechanical Engineering, LBS Institute of Technology for Women, Poojappura.
Email: sajanjerome@lbsitw.ac.in, ORCID id: https://orcid.org/0000-0003-4231-5265
[6]Professor, Mechanical Engineering, LBS Institute of Technology for Women.
Email: anilkumar@lbsitw.ac.in, ORCID id: https://orcid.org/0000-0001-5185-9489

ABSTRACT

There is a growing use of artificial intelligence in healthcare marketing in the form of chatbots, predictive outreach, recommender systems, segmentation engines, digital front-door workflows, and generative content systems. Although such technologies are meant to enhance personalisation and efficiency in operations, healthcare is not like an average market of consumers due to its vulnerable users, highly sensitive information, and the decision-making process, which can affect access, trust, and behaviour related to health. The present paper will carry out a research study of the ethical AI implementation in healthcare marketing through a management lens, based on secondary data. The discussion uses scholarly works published by peers, policy documents, and regulatory guidelines, and enforcement practices of healthcare organisations to explore the ways in which AI can be used by healthcare organisations in a way that is patient-centric, ethically justifiable, and operationally feasible. The literature review incorporates the literature on patient engagement, interaction via omnichannels, AI ethics, privacy governance, fairness, transparency, and interactions in digital health. The results indicate that predictive models are not always the root cause of the most severe ethical breaches in healthcare marketing, but the entire marketing-technology stack: web trackers, third-party advertising techniques, SDKs, and vendor ecosystems. The conceptual model proposed in the paper suggests a management-based connection between the computer-science capabilities and the marketing capabilities to the patient-centric strategies, mediated by the lifecycle governance principles which are accountability, fair evaluation of consent, transparency, and risk-monitoring. Findings indicate that organisations ought to view ethical AI as an organisational capability, other than a compliance exercise. The factors that are needed to maintain trust and facilitate responsible innovation are strong governance, privacy-safe measurement, well-design consent, human-based oversight, and vendor accountability. The paper ends with practical recommendations to managers and finding limitations and directions of future research**.**

Keywords: Ethical AI; healthcare marketing; patient-centricity; responsible AI; privacy governance; algorithmic fairness; explainability; digital health; generative AI; healthcare management.

How to cite this article: Resmi R, Kruthika CG, Nafeza E, Gayathri A, Jerome S, Anilkumar EN. Ethical AI Deployment in Healthcare Marketing: A Management Perspective on Computer Science-Enabled Patient-Centric Strategies. Int J Drug Deliv Technol. 2026;16(11s): 297-305; DOI: 10.25258/ijddt.16.11s.28.

1. INTRODUCTION

A structural change in healthcare marketing has been a transition to ongoing digitally mediated communication of the brand throughout the patient journey instead of periodic brand communication. Online searches, hospital websites, patient portals, mobile applications, teleconsultation

platforms, online communities, and CRM systems have become the new means through which patients find providers, compare or contrast treatment alternatives, make appointments, process health information, and have continued relationships with healthcare organisations. In this context, marketing is no longer a purely promotional activity. It has become more and more overlapping with patient experience design, service access, care navigation, and relationship management.

This change is magnified by artificial intelligence, which allows the process of personalisation and automation on a large scale. Needs of the services, audience segmentation, no-show prediction, and promotion of content or booking directions can be identified using machine learning models. NLP serves chatbots, semi-Automatic triage-related navigation, multilingual interaction and sentiment analysis. Real time Generative AI can generate health education summaries, reminders copy, campaign variation, and dialogue scripts. Managerially, such systems will be efficient, responsive, and have an improved conversion of digital contact to service uptake.

Nonetheless, the field of healthcare marketing is different ethically as compared to the traditional commercial marketing. Patients tend to engage in healthcare systems when they are vulnerable, uncertain and asymmetrical in knowledge. In most jurisdictions, health information is considered to be sensitive or special-category data. A decision in marketing may not only have an impact on consumer preference: it may also impact access to care, treatment understanding, compliance, stigma exposure, and trust in the providers and digital systems. Hence, behaviour that might appear normal in a retail store setting, like profiling aggressively, opaque optimisation, or blanket third-party tracking, turns into an issue in a healthcare setting.

These risks have been witnessed in the recent cases. Public disclosure and regulatory intervention have revealed that certain digital health and healthcare organisations have sent sensitive health related signals to advertising and analytics services in pixels, cookies, SDKs, or similar tools. The occurrence of such instances shows when ethical harm arises as a result of the larger data and vendor ecosystem as opposed to an individual algorithmic model. In health, however, the application of AI ethically cannot be achieved merely based on technical performance; it demands managerial discretion, governance structure and patient-centred responsibility.

The research question that this paper will answer is the following: how can healthcare organisations engage AI in marketing in a manner that is provably patient-centric, ethical, and strategically sustainable? The study is answered through the secondary-data design and is divided into a traditional research style with literature review, methodology, analysis, findings, limitations and conclusion. The article makes a contribution by bringing together various pieces of evidence to an operational management framework of ethical AI implementation in healthcare marketing.

2. LITERATURE REVIEW

2.1 Patient-centricity and healthcare value creation

Patient-centricity has gained the centre stage of healthcare management discourse in the last 10 years. As opposed to institution-centred models, patient-centricity focuses on the choice, understanding, agency, and experience of patients throughout the entire care process. The current body of research studying the concept of patient engagement and service co-creation reveals that, in healthcare, value is not passively provided by healthcare providers, but rather is co-created through the quality of communication, the design of access, informed-participation, and relational trust. This literature is particularly applicable to the context of marketing in healthcare by ensuring that the marketing strategies that are ethically based should lead to less friction, better understanding and decision-making, as opposed to simply maximising response or conversion.

Researchers of patient engagement suggest that meaningful engagement cannot be delivered in the form of customized messages only. It entails organisational processes that encourage the patient to browse, to comprehend trade-offs and to take part in care-related decisions. Patient-centric marketing needs to be measured in this way in terms of contribution to informed access and continuity, not only in terms of consumer measures. A hospital web site which assists a patient comprehending specialty pathways or reserving the appropriate service effectively can be more patient-centred than a massively customised campaign which results in clicks whilst not ensuring that care is more relevant.

2.2 Omnichannel healthcare interaction

The other significant literary source is that of omnichannel healthcare interaction. The healthcare user now experiences provider organisations on both digital and physical platforms, and their continuity expectations are informed by their experience in banking, retail, travel and consumer technology. The studies in this field indicate that well-coordinated multi-channel interaction can enhance the access, communication and retention under the condition of good infrastructure, analytics and personnel capacity. Digital front doors are becoming increasingly popular in healthcare organisations encompassing search, web, chat, scheduling, reminders, and feedback loops.

However, there also exist significant tensions in the literature of the omnichannel. Discontinuous systems may also cause duplication, omission, and confusion, particularly among the aged, rural, or low digital literacy patients. These tensions may be magnified when AI is overlaid on the omnichannel systems. As one example, predictive outreach could help enhance access in certain groups but leave others undetected due to targeting by proxy or due to poor data coverage. Thus, the omnichannel point of view reinforces the thesis that AI use in healthcare marketing should be measured not only based on the standards of operational throughput but also with regard to equity and patient experience.

2.3 Responsible AI and ethics frameworks

A third key literature flow is that of responsible AI schemes and ethics. The documents on global governance formulated by the World Health Organization, OECD, UNESCO, the European Union, NIST and national policy agencies focus on common themes of respect to human agency, fairness and non-discrimination, transparency, accountability, privacy and data governance, safety and societal well being. The usefulness of these frameworks lies in the fact that they offer normative reference points to assess the AI system and underline the impossibility of limiting trustful AI to technical accuracy.

With regards to health in particular, WHO guidance claims that AI must safeguard autonomy, enhance well-being and safety, guarantee transparency, responsibility and accountability, inclusiveness and equity, and stay responsive and sustainable. The credible EU AI framework also emphasizes on human control, technical strength, privacy, fairness, transparency, and responsibility. The principles are converted into an operational lifecycle of govern, map, measure, and manage by the AI Risk Management Framework by NIST. This is especially helpful to the managers since it transforms ethics into an abstract principle to a repeatable process of governing.

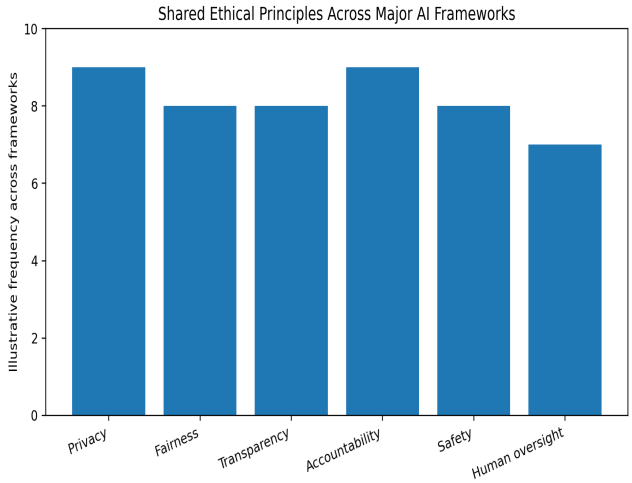

Figure 1. Shared ethical principles across major AI frameworks (author-created).

Meanwhile, as observed in the critical literature, most of the AI ethics systems are still high-level and principle-based. Their under-specifications on what organisations must do with the actual operational dilemmas like ad-tech vendor dependencies, dark-pattern consent, performance trade-offs across patient groups, or the governance of generative content tools, are common. This disconnect is particularly noticeable in healthcare marketing which has technology stack consisting of CRM systems, attribution tools, analytics tags, campaign automation, tracking scripts, and third-party media.

2.4 Privacy, consent, and health data governance

The two topics dominate the literature of legal and scholarly studies of healthcare AI; they are privacy and consent. The information concerning health-related issues is generally better protected under the substantial regulatory systems since such data can easily expose intimate personal conditions, vulnerability, or unpredictable risks. According to legal commentary on the data protection law, it is emphasized that behavioural cues related to health might be sensitive, even when compiled outside of the conventional clinical records. This has significant implications on healthcare marketing since browsing behavior, the desire to make an appointment, symptom search, and portal behavior can all serve as health status proxies.

The studies of mHealth, consent interfaces, and transparency are revealing that users will typically have a hard time knowing how their data will travel through complicated digital ecosystems. Formal consent can be there without substantial understanding. Practically, consent obtained through dark patterns, incomprehensive privacy policy, baked-in permissions, and silent third-party transfers may be legally weak and morally wimpy. In the case of healthcare managers, this should mean that consent has to be developed as a customer experience element and not as a legal check-box exercise.

The fact that there are parental enforcement measures and investigations confirms this. A number of digital health and health-related applications have been questioned regarding sharing sensitive health information with advertising services despite the privacy claims that they provided. The cases indicate that secondary use and concealed disclosure is often the ethical issue, not just in model training. Thus, the healthcare marketing governance should cover strict consideration of tag management, vendor checks, reduction of telemetry, and mapping of data flows.

2.5 Trust, explainability, and patient acceptance

Another significant driver in the adoption of AI has been the patient trust. Empirical research on the attitudes of the population and patients regarding AI integration in health systems has highlighted that perceptions of usefulness and

trust in organisations using the data in a manner that would not lead to discriminatory effects and retaining meaningful human control are largely determinants of acceptance. There is relational and institutional trust. Patients can easily accept automation in case they believe that a provider is accountable and open, yet the mistrust builds up very fast in case AI turns out to be manipulative, invasive, or unresponsible.

This is something to do with explainability. The marketing systems are in most cases targeting or sequencing content in a manner inaccessible to the patients. When a person is bombarded numerous times with information about a condition, a service line or a treatment pathway, without knowing the reasons behind it, the experience can be invasive or forceful. This issue is more acute in such sensitive areas as mental health, reproductive health, treatment of addiction, or support of chronic illnesses. Explainability in healthcare marketing, therefore, cannot be understood as just technical model interpretability, but also intelligible communication regarding the use of the data, logic of targeting, and possible options.

2.6 Fairness and bias in healthcare-related algorithms

Another important dimension is added to bias and fairness research. The literature on the bias in healthcare algorithms is well-known since it demonstrates that proxy variables may reproduce structural inequalities when optimisation objectives are not related to the real need or welfare. This lesson has a direct correlation with healthcare marketing. When systems are optimised to serve a purpose of click-through/conversion/revenue potential/ historic utilisation, they can be under serving groups with less history of access, less digital literacy, or less strong engagement footprints. That is, optimisation may take inheritance of inequality and redefine it as efficiency.

Justice in healthcare marketing thus involves widened outcome measures. Instead of posing the question whether a certain group of people clicks more than the other, the managers need to pose the question whether the AI-facilitated outreach will lead to a more equitable access, understanding, retention, and navigation between groups. This is the change of the limited marketing metrics into the patient outcome proxies, which is the core of a patient-centric management perspective.

2.7 Generative AI and communication risk

The fast development of generative AI gives new possibilities and new threats. Big language models could be used in multilingual support, content personalisation, summarisation, and cheap generation of messages. Such affordances are appealing in healthcare communication particularly to organisations that are being pressurised to scale patient engagement effectively. However, generative AI is also associated with hallucinated assertions, hypersimplification, tone shift, implicit bias, loss of privacy, and lack of source control.

Generative tools in the marketing of healthcare services should also be regulated closer to safety-relevant communication systems than to traditional copywriting assistants. Messages have the ability to affect the decisions of service, risk perception, medication comprehension or disclosure intentions. The fact that the literature on language-model harms contributes to the argument that human review, source grounding, prompt controls, audit logs, and policy restrictions should be introduced in the health communication workflows.

2.8 Literature gap

This is demonstrated by the literature in general: ethical AI implementation in healthcare marketing is at the crossroads of patient-centricity, digital service management, privacy control, fairness, and responsible innovation. However, a clear gap remains. Several of the literature discusses either general AI ethics in the healthcare field or operational marketing practice, but there is less literature that links them using a management framework tailored to healthcare marketing. The gap that is selected in this paper is the synthesis of these streams in secondary-data, which is aimed at the managerial implementation.

3. METHODOLOGY

The research design used in this paper is a qualitative secondary-data design. The research based on secondary data is suitable since the topic will be a synthesis of several other evidence areas: peer-reviewed studies, ethics guidelines, regulatory reports, and case studies of the population. The task is not to statistically test a single causal model but to come up with an analytically based and managerial application of ethical AI implementation in healthcare marketing.

There are four categories of the source base. To begin with, the peer-reviewed literature was analyzed regarding the literature connected to patient engagement, omnichannel healthcare interaction, healthcare marketing, chatbot ethics, patient trust in AI, fairness, privacy, and generative AI risk. Second, normative requirements and governance were identified with the help of policy and standards documents. These consisted of large-scale international and national AI ethics, data governance, and digital regulation of healthcare-related issues. Third, enforcement practices and regulatory directions on the part of the government were revisited so as to record actual trends of failure, as far as marketing and health data is concerned. Fourth, the selection of industry and institutional material was made to put into context the present-day practices of healthcare marketing and managerial realities.

Relevance, credibility and recency were used as guides when selecting the sources. Preference was given to the materials published since 2016, but seminal earlier works were provided where needed to provide the theoretical background. They were analyzed with the help of thematic synthesis. Deductive first coding categories were formed out of common principles of ethics including privacy, consent, fairness, transparency, accountability, safety, and human oversight. Other inductive coding included healthcare marketing practices specific patterns including tracking technologies, vulnerable population sectors,

vendor risk, conversion optimisation and content governance.

Since this is a secondary and qualitative study, it has the purpose of analytical generalisation as opposed to universal statistical inference. The result is a theoretical model and a list of managerial suggestions based on thereathedeviations of evidence of various types.

4. CONCEPTUAL FRAMEWORK

The paper suggests a integrative model where the ethical AI implementation is addressed as an organisational management ability. This model has three domains of input interacting. The first one is computer-science capability, such as data engineering, machine learning, NLP, generative AI, cybersecurity, and interoperability. The second area is marketing capability and it encompasses segmentation, CRM, campaign orchestration, content operations, design of digital experience, and analytics. The third sphere is the situation of governance, such as regulation, clinical norms, privacy expectations, risk handling, and the oversight of vendor ecosystem.

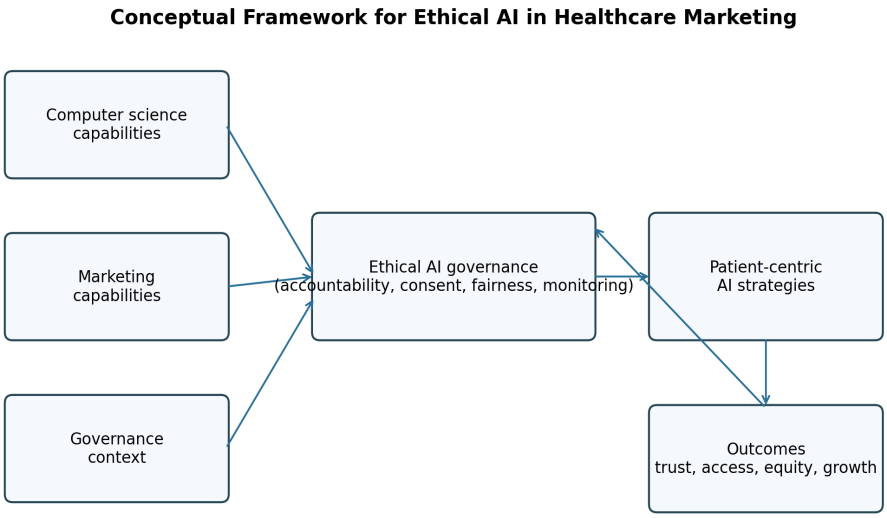

Figure 2. Integrative conceptual framework for ethical AI governance in healthcare marketing (author-created).

Collectively, these domains facilitate AI-assisted patient-centred practices, e.g. digital navigation, personalised education, proactive outreach, conversation assistance, service recovery communication and continuity support. The efficacy and validity of these strategies however are tempered by governance functions: accountability, structures, lawful and meaningful consent, transparency, fairness assessment, human review, minimisation of data and post deployment auditing. In the case of weak governance functions, AI capabilities can result in short-term growth or efficiency and long-term sabotage of trust, equity, or safety.

The structure thus connects the ethical governance with strategic performance. Trust, fair access, patient satisfaction, reputational strength, and sustainable digital development are not perceived as independent objectives but are interdependent.

5. ANALYSIS AND RESULTS

5.1 Clustered uses of AI in healthcare marketing

The secondary evidence allows stating that the application of AI in healthcare marketing can be grouped into five larger clusters.

The former refers to conversational assistance. Hospitals, clinics, insurers and digital health platforms are increasingly using AI facilitated chatbots and virtual assistants to answer general queries, redirect users to service lines, help make appointments and answer general information. These instruments can reduce friction and increase responsiveness especially not during the working hours. Their ethical risks include being exposed to privacy, being misinformed, inappropriate confidence and failure to deal with vulnerable users.

The second one is predictive outreach. Organisations use AI to recognize the risk of no-show, most likely disengagement, or care-gap patterns or service propensity. Such outputs would be beneficial in follow-up messages and continuity campaigns and reminders. Well managed such systems are able to enhance access and reduce care disruptions. They may be biased in their construction, or may overrepresent one group or may not represent underserved groups with less robust data histories.

The third category is individualised learning and content modification. It may be applied to personalize learning experiences through the use of AI tools that combine according to language, literacy level, previous interaction, and mediator. This will be capable of supporting patient

comprehension and improve navigation. However, personalisation of the content is not ethical when it augers risks, simplifies the choice architecture or taps into perceived vulnerability to deliver emotionally manipulative content.

The fourth group is the segmentation and orchestration. AI-driven CRM systems are able to prioritize audiences, sequences of messages, and optimise campaigns. This is likely to be the closest to the normal commercial marketing. In healthcare, however, this optimisation is cautiously undertaken as the audience that is considered valuable is not always the most historically-converted, but it could simply be a tendency of privileged admission. Pure conversion logic may result in patient-centricity being undermined.

The fifth group is measurement and attribution. Here, the majority of the gravest ethical problems are raised. The marketing teams often use analytics scripts, cookies, advertising pixels, SDKs and third-party tools to measure the success of a campaign. Being implemented into the healthcare setting, these tools can unwillingly elevate the mundane telemetry to the healthcare-related information streams. As a result, the ethical risk may be the most serious in either the invisible elements of AI or the infrastructure.

5.2 Key ethical risk patterns

The five large patterns of ethical risk are analyzed.

The latter happens to be the invisible secondary utilization of health-related data. Normal online communication may include the intention of appointing, interest in the symptoms, research into the route of treatment, or an association with the provider, which may manifest in a marketing exercise. Once they are disseminated to the third parties, such data may be employed to characterize the patients in a manner they do not anticipate or even accept.

The second one is the lack of strength in consent. Healthcare organisations can provide privacy notices or consent banners technically but the real experience of using the service does not always favour informed choice. Consent is prone to undue length of notices, prior selection of default on notice, inconvenience of withdrawal and conception, and aptitude of purposes. In this case, the ethical aspect of AI-based personalisation is tainted.

The third is proxy bias and unreasonable optimisation. The systems may also give precedence to the better served populations in the event they access it through revenue, past interactions or past use of the service. Systematically less connected, less digitally assured or previously suppressed demand groups can be deprioritised. It is a strategic problem and an ethical problem since it can lead to a rise in inequalities and weaken institutions.

The fourth theme is the theme of transparency. The patients do not even have much knowledge of why they are being targeted, redirected or even sold a particular chain of message. Even within teams may not be so apparent with the logic of the campaign being disseminated by multiple software vendors. This makes the problem of responsibility more difficult and weakens the authority to justify, question or correct judgments.

The fifth style is the generative communication risk. The duplicate generated by AI can be utilized to scale up operations, but, unless closely edited it may generate false or mis-leading statements, especially in controlled or clinically vital fields. This is not just a branding risk in healthcare; it can affect decisions and trust among the patients.

5.3 Management interpretation of results

As a management matter, the findings indicate that the implementation of ethical AI cannot be transferred to technical units or compliance managers only. It involves coordination of the enterprise marketing, patient experience, legal, information security, data governance, and clinical leadership. The major managerial issue is not whether to implement AI, but how to structure organisational controls that would match AI application to patient-centric objectives.

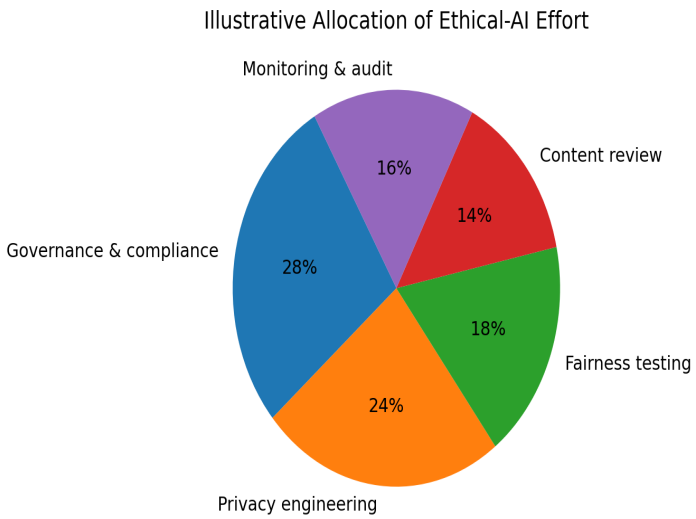

Figure 3. Typical allocation of ethical-AI effort in a healthcare marketing programme (author-created).

There is also evidence that healthcare organisations must not regard personalisation as the main indicator of sophistication. A very individualised system based on opaque sharing of data, problematic targeting, or biased optimisation can be less patient centred than a less complex system which is built around transparency, accessibility, and navigation assistance. Therefore, matured restraint may be an expression of being mature and not weak.

6. FINDINGS

Five major findings are based on the analysis.

Finding 1: Ethical risk in healthcare marketing is often brought about in the marketing-technology ecosystem around it and not in isolated AI models. The areas of central governance are therefore tracking architecture, vendor integrations and measurement choices.

Finding 2: Patient-centricity cannot be synonymous with personalisation only. A patient centred approach is one that promotes autonomy, clarity, accessibility and continuity without infringing privacy and meaningful choice.

Conclusion of the finding 3: Equity Conversion-centred optimisation has the capacity to recreate inequity when training data and campaign logic are based on historical patterns of access. An assessment of fairness should then be on the quality of access and engagement among groups as opposed to the overall campaign performance.

Finding 4: Trust is a strategic resource in AI-driven healthcare marketing. Not only is transparency, consent design, human oversight, and responsiveness to complaints imperative in compliance, but also long-term legitimacy.

Finding 5: Generative AI makes healthcare communication a more risky field requiring content control, review mechanisms and source management like other healthcare functions that are of critical quality.

7. DISCUSSION

The discussion supports the idea that the implementation of ethical AI in healthcare marketing is not a technical undertaking. It is a management skill that combines technology design, managerial control, legal consciousness and patient experience mind. This approach changes the strategic answer of the question, How can AI improve marketing performance to the question, How can AI improve patient-centred interaction without violating rights, trust, or equity?

An important implication is that AI metrics should have patient-centric outcomes that have been aligned by managers. Conventional measurements including click-through rate, cost per acquisition and conversion are still applicable although they are not adequate. Complementary metrics need to be appointment completion, message understanding, opt-out, complaint rates, benefit distribution within groups, and trust metrics. The measures will assist in avoiding systems that are optimised to attain small business results at the cost of large organisational legitimacy.

The other implication is on the issue of vendor governance. The stacks of modern healthcare marketing are frequently put together using a number of platforms, many of which deploy black-box processes or secondary processors. Managers are therefore advised to enforce inventory management, contractual limitations, audit privileges and change management to track tools, APIs, SDKs, and services of the models. It is arguably acceptable in sensitive service lines to use a default no health data for advertising position, than to gauge tolerable leakage levels.

It should also focus on the patient experience lens. Ethical AI must be perceived as less spectacular technology but rather less friction, provide clarity, be written in understandable language, have a respectful timing and be informed choice. When a patient is able to interpret why they received a message, to manage data use and find the correct service, he/she is more likely to feel that AI is beneficial and not abusive.

8. LIMITATIONS

There are a number of limitations in this paper. First, it is founded on the secondary data and analysis is relying on published evidence, publicly available structures, and observable cases. Enforcement activities by the public can disproportionately highlight the existence of extreme failures, and may fail to record less conspicuous internal successes or near-misses. Second, the regulatory environments change fast, especially in the context of AI and digital health, and the interpretation of governance should not be considered as a legal consultation. Third, the research summarises the literature of various jurisdictions and local law might vary considerably. Fourth, since the design is qualitative and theoretical, the framework has not been empirically verified with the help of interviews, surveys, or case comparison.

These weaknesses do not render the study worthless, but rather, a primary research is necessary in future. The practical usability of the proposed framework may be tested by the interviews with marketing leaders, compliance professionals, clinicians, and patients. The model might also be optimized with the help of comparative case studies among different hospitals, insurers, and digital health platforms.

9. CONCLUSION

The ethical use of AI in healthcare marketing can be most appropriately considered a technology-governance-patient-experience issue that is a management challenge. AI can assist healthcare organisations to become more responsive, navigating, continuity, and quality of communication. But with poorly governed governance, the same abilities can also be used to do intrusive tracking, unfair optimisation, opaque targeting, and insecure automated content.

The main point of this paper is that patient-centric healthcare marketing involves more than the sophisticated analytics or the personalised campaigns. It demands a rigorous governance design that ensures that AI implementation is in line with privacy, fairness, transparency, consent, and human accountability. Ethical AI should not be seen as a compliance tool by managers, but as an organisational capability that has a direct relationship with trust and long-term strategic results.

An accountable way to go also entails accountability among the executive, privacy-appropriate measurement, relevant consent design, fairness monitoring, human examination of delicate material, and a high degree of oversight over the vendors. With such practices, the organisations will be in a better position to leverage on AI without jeopardizing the trust of the patients and the authenticity of the medical institutions.

REFERENCES

1. Advertising Standards Council of India. (n.d.). The ASCI code.
2. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the ACM Conference on Fairness, Accountability, and Transparency.
3. Centers for Disease Control and Prevention. (2024). Health equity and ethical considerations in using artificial intelligence in health care and public health.
4. Department of Health and Social Care. (n.d.). Digital and data-driven health and care technology: Code of conduct and good practice guidance.
5. Dwork, C. (2008). Differential privacy: A survey of results.
6. European Commission. (2020). Assessment list for trustworthy artificial intelligence (ALTAI).
7. European Commission High-Level Expert Group on AI. (2019). Ethics guidelines for trustworthy AI.
8. European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation).
9. European Union. (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act).
10. Food and Drug Administration. (2024). Using social media: Guidance and notices.
11. Federal Trade Commission. (2021). FTC finalizes order with Flo Health.
12. Federal Trade Commission. (2023a). FTC enforcement action to bar GoodRx from sharing consumers' sensitive health info for advertising.
13. Federal Trade Commission. (2023b). FTC gives final approval to order banning BetterHelp from sharing sensitive health data for advertising.
14. Federal Trade Commission. (2023c). Easy Healthcare Corporation, U.S. v. case materials.
15. Federal Trade Commission. (2024). Monument, Inc. settlement materials.
16. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. Communications of the ACM, 64(12), 86–92.
17. Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. Minds and Machines, 30, 99–120.
18. HHS Office for Civil Rights. (n.d.-a). Marketing (HIPAA Privacy Rule guidance).
19. HHS Office for Civil Rights. (n.d.-b). Use of online tracking technologies by HIPAA covered entities and business associates.
20. India Code. (2023). The Digital Personal Data Protection Act, 2023.
21. International Organization for Standardization. (2023a). ISO/IEC 23894: Artificial intelligence—Guidance on risk management.
22. International Organization for Standardization. (2023b). ISO/IEC 42001: Artificial intelligence—Management system requirements.
23. JAMA Network Open. (2024). Patients' trust in health systems to use artificial intelligence.
24. McKinsey & Company. (2023). Marketing in healthcare: Improving the consumer experience.
25. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting.
26. National Health Authority. (n.d.). Ayushman Bharat Digital Mission.
27. National Health Authority. (n.d.). Health Information Exchange and Consent Manager brochure.
28. National Institute of Standards and Technology. (2023a). Artificial Intelligence Risk Management Framework (AI RMF 1.0).
29. National Institute of Standards and Technology. (2023b). NIST AI RMF Playbook.
30. NHS England Transformation Directorate. (2025). Artificial intelligence: Information governance guidance.
31. NITI Aayog. (2021). Towards responsible AI for all: Approach document for India.
32. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453.
33. OECD. (2019). Recommendation of the Council on Artificial Intelligence.
34. OECD. (n.d.). AI Incidents and Hazards Monitor.
35. Rieke, N., Hancox, J., Li, W., et al. (2020). The future of digital health with federated learning. npj Digital Medicine, 3, 119.
36. Sundaramurthi, T., Mathews, S., Bermudez, E., & Mahajan, S. M. (2024). Patient engagement and co-creation in healthcare services: A scoping

review. Patient Experience Journal, 11(3), 215–245.
37. The Markup. (2022). Facebook is receiving sensitive medical information from hospital websites.
38. UNESCO. (2021). Recommendation on the ethics of artificial intelligence.
39. United States OSTP / White House. (2022). Blueprint for an AI Bill of Rights.
40. Weidinger, L., et al. (2021). Ethical and social risks of harm from language models.
41. World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance.