

# Beyond consent: Reconstructing ethical justification in medical adaptive machine learning systems

Keiichiro Yamamoto<sup>1,\*</sup>, Makoto Udagawa<sup>2</sup>, Eisuke Nakazawa<sup>3</sup>

<sup>1</sup> Department of Clinical Research Management, Center for Clinical Sciences, Japan Institute for Health Security, Tokyo, Japan;

<sup>2</sup> Section of Bioethics, Department of Clinical Research Support, National Center of Neurology and Psychiatry, Tokyo, Japan;

<sup>3</sup> Department of Biomedical Ethics, Faculty of Medicine, The University of Tokyo, Tokyo, Japan.

**Abstract:** Medical Adaptive Machine Learning Systems (MAMLS) that continuously update their models using clinical data blur the conventional boundary between therapy and research, prompting the argument that their use should be classified as research and governed by informed consent requirements. Although informed consent remains normatively and legally important, this paper contends that consent-centered ethics faces two structural limitations in the context of MAMLS. First, the irreversibility inherent in deep learning models substantially undermines withdrawability—an important ancillary right of consent—thereby suggesting that consent may be transformed from an instrument of ongoing self-determination into a form of delegation to institutions. Second, the problem of data representativeness and bias shifts the unit of ethical analysis from the individual to the population, creating an "autonomy dilemma" in which respect for individual consent can paradoxically undermine the protection of autonomy at the collective level. Under these conditions, ethical justification must be complemented by, and in some contexts repositioned toward, public trust in institutions. The paper concludes that the ethical challenges surrounding MAMLS cannot be adequately addressed within the framework of research ethics alone, but must instead be taken up within the broader framework of public health ethics, with particular attention to transparency, accountability, and participatory governance.

**Keywords:** medical artificial intelligence, informed consent, public trust, research ethics, public health ethics

## 1. Introduction

In recent years, there has been growing interest in artificial intelligence (AI) in healthcare—particularly Medical Adaptive Machine Learning Systems (MAMLS)—with increasing attention to their potential clinical deployment (1,2). Unlike conventional medical AI systems that rely on fixed, pre-trained models, MAMLS continuously update their models by incorporating patient data generated in clinical settings after initial deployment (1,3). In this sense, MAMLS are not merely tools for diagnosis or treatment support; they are technologies in which the act of use itself encompasses a process of knowledge generation. Since the clinical environment thus doubles as a research environment, the distinction between "therapy" and "research"—a distinction the Belmont Report (1979) deemed essential to maintain—becomes increasingly blurred, both institutionally and conceptually (4,5).

On this point, Sparrow and colleagues argue that, given the continual-learning nature of MAMLS,

their use should be classified as research (6). Their argument can be summarized as follows: since MAMLS continuously update their models by incorporating patient data in clinical settings, each patient's data contributes to a process that produces generalizable knowledge affecting future patients, thereby structurally satisfying the Belmont Report's definition of "research" (4). Furthermore, patient data is used not only to optimize individual treatment but also to improve the model, meaning that patients are not necessarily involved solely for their own benefit. Given that manufacturers have financial interests in system improvement, they argue that continual learning without treating patients as research subjects gives rise to a risk of moral hazard. Consequently, they conclude that IRB oversight and prospective written opt-in informed consent are ethically required (6).

This argument is consistent with the mainstream tradition of research ethics shaped by the Belmont Report. However, the consent-centered regulatory vision appears to be in tension with actual institutional

trends. In Japan, the "three-yearly review" of the Act on the Protection of Personal Information has examined legislative designs that would permit AI development and statistical use without individual consent in certain circumstances (7). The Personal Information Protection Commission's reform policy has raised as an issue the appropriate role of individual consent in facilitating AI development understood as a form of statistical processing, and the public consultation process generated numerous responses concerning "data utilization that does not require individual consent" (7,8). Among these responses, some characterized the relaxation of consent requirements as a shift from *ex ante* regulation to *ex post* governance (8,9). These reform discussions are still ongoing, and the final legislative design has not yet been determined. Although these developments are specific to Japan, they illustrate a broader international challenge in medical AI governance: how to reconcile large-scale, socially valuable data use with the continuing normative importance of individual consent, accountability, and public oversight. They therefore provide a useful institutional example of the wider shift from *ex ante* consent-centered regulation toward *ex post* governance and trust-based accountability.

What this institutional context reveals is a marked gap between normative principles and practical realities: while consent is positioned as central to AI development from the perspective of research ethics, it is simultaneously recognized as being in tension with large-scale data utilization in practice (10-12). Particularly in MAMLS, it is necessary to continuously acquire broad and minimally biased data in order to adequately cover the relevant data distribution and thereby ensure model performance and safety (13). In other words, data-related bias must be minimized as far as possible. This creates a marked divergence between norm and practice: on the one hand, stronger consent requirements are advocated from the standpoint of research ethics, while on the other hand, consent requirements are being relaxed in the name of AI utilization policy.

This paper seeks to reframe this tension not merely as a clash of policy choices, but as a manifestation of the structural limits of consent as conventionally understood in research ethics. First, it argues that MAMLS inherently involve irreversibility, which substantially undermines withdrawability as an important ancillary right of consent. Second, it contends that problems of data representativeness and bias shift the unit of ethical analysis from the individual to the population. Third, it argues that, as a consequence, the focus of ethical justification should be repositioned from individual consent to public trust in institutions. On this basis, the paper concludes that the ethical challenges posed by MAMLS cannot be adequately addressed within the framework of research ethics

alone, but must also be taken up within the broader framework of public health ethics.

## 2. Irreversibility and the breakdown of consent

To understand the characteristics of MAMLS from the perspective of research ethics, it is first necessary to examine their underlying technical features. Many contemporary medical AI systems are based on deep learning models, including convolutional neural networks (14,15). In such models, training data is not preserved as discrete records, but is incorporated into the model through changes to its internal parameters (16). Because data is embedded throughout the model's internal structure, it is difficult to determine how any particular data point has influenced the model.

Under such a structure, selectively deleting data once incorporated into training and restoring the model to a state in which that data had not been incorporated is technically and operationally difficult under current conditions (17). While methods such as retraining and machine unlearning are theoretically conceivable and increasingly explored, they may require substantial computational resources, time, and system-level reconstruction in practice (18). Retroactively removing only the contribution of specific data may therefore be difficult to implement reliably in deployed MAMLS. Data utilization in MAMLS thus appears to take on a practically irreversible character under realistic technical and operational conditions. Sparrow and colleagues note that MAMLS carry the risk of "catastrophic forgetting"—a phenomenon in which the model overwrites previously acquired information during the update process (6)—but the implications of such practical irreversibility for the concept of consent itself do not appear to have been sufficiently addressed.

This irreversibility stands in tension with the philosophical premises underlying informed consent. In both research ethics and clinical ethics, informed consent has generally been understood to include withdrawability as an important ancillary right, alongside informedness, understanding, and voluntariness (19). Withdrawability is a crucial condition for understanding consent not as a one-time act of permission, but as a temporally extended process of self-determination (20,21). Within research ethics in particular, ensuring the voluntariness and withdrawability of consent to participation is a central implication of the principle of respect for persons and an important safeguard against unjust exploitation of research participants (22,23).

In MAMLS, however, this premise is significantly destabilized. First, since data is integrated into the model in ways that may be difficult to reverse, even if the formal right to withdraw consent is recognized, it is difficult to give substantive effect to that withdrawal. Second, given the nature of continuous learning, the

scope and impact of data use change over time, and the full picture cannot be grasped in advance. Under these conditions, consent necessarily becomes something given comprehensively in relation to indeterminate future uses—less a reversible choice than a commitment whose practical consequences may not be fully undone.

Conventionally, consent has been understood as a means by which individuals exercise ongoing control over matters concerning themselves (20,21). In MAMLS, however, such control is severely constrained by technical and operational limitations. Consequently, consent takes on a character closer to "delegation"—an act of entrusting discretion regarding future uses and their consequences to institutions and professionals (24). The question, therefore, is under what conditions such delegation can be justified in a setting that presupposes data use that is difficult to reverse in practice. This, in turn, opens onto broader questions about how ethical justification is to be secured within institutional and social arrangements.

### 3. From individuals to populations

The characteristics of MAMLS examined above are not limited to irreversibility: through the ways in which data is collected and used, they also potentially reconfigure the very unit of ethical analysis. The focus of the problem thus appears to shift from the choices of individual patients to the nature of the populations constituted through data.

In MAMLS, the quantity and diversity of training data are critical determinants of model performance and safety. Particularly in the medical domain, data reflecting diverse attributes—such as age, sex, genetic background, and socioeconomic status—are required (13). However, opt-in data collection necessarily depends on participants' choices and therefore tends to result in certain groups being over- or underrepresented—the well-known problem of selection and sampling bias (1). This problem has already been observed in real-world systems. In dermatological image-recognition AI, for example, diagnostic accuracy for patients with darker skin tones is significantly lower owing to the overrepresentation of White patients in the training data (25,26). Consequently, biases in training data can reduce the model's predictive accuracy for particular groups.

Such data bias is not merely a technical problem. Through reduced diagnostic accuracy and increased misdiagnoses for certain patient groups, it gives rise to substantive disadvantage and thereby raises concerns about both scientific and ethical validity. Sparrow and colleagues also recognize the problem of bias: they argue that, in MAMLS that employ collective learning, models may not necessarily improve—and may even worsen—for particular subpopulations, thereby raising the ethical concern that such research subjects may

be treated as "mere means" (6). However, they seek to address this problem within the framework of risk-benefit assessment in research ethics. What this paper suggests, by contrast, is that the problem may involve a more fundamental transformation—one that shifts the very unit of ethical analysis from the individual to the population—and therefore cannot be adequately contained within such a framework.

Here, an ethical tension becomes manifest. The more thoroughly individual autonomy is respected and consent-based data provision is pursued, the more likely it is that data bias will arise, potentially leading to a decline in model performance. Conversely, efforts to ensure the comprehensiveness and representativeness of data tend to reduce the scope for individual choice and weaken the substantive meaning of consent (27). This situation represents what might be called an "autonomy dilemma": efforts within the traditional research ethics framework to secure respect for persons through consent-centered mechanisms may paradoxically undermine the effective protection of data providers across the population, including their autonomy-related interests.

Under this dilemma, the very unit of ethical analysis is reconfigured. Conventional research ethics has been built primarily around individual patients and research participants, with a primary focus on their rights and interests (4,22). In MAMLS, however, individual data items have no meaning in isolation; they fulfill their function only as part of an aggregate. The focus of the problem thus shifts from the individual question of "who consented" to the structural question of "what kind of population is constituted through data". This shift demands a change in ethical framework: from protecting individual research participants to protecting the broader population that provides clinical data, with greater emphasis on collective interests, distributional fairness, and the distribution of risks—thereby suggesting the need for a shift toward public health ethics (28,29).

### 4. From consent to trust

As we have seen, irreversibility and collectivity fundamentally undermine the conventional consent-centered ethical framework in MAMLS. What is at issue is not merely that consent procedures are inadequate; rather, the question is whether consent itself can function as an adequate principle of justification in this context. This argument should not be read as denying the normative or legal importance of consent. Rather, consent remains an important expression of respect for persons and may retain legal significance even when it cannot, by itself, bear the full justificatory burden of MAMLS governance.

One source of this difficulty lies in the limits of understanding. The internal structure of deep learning

models is highly complex, and it is difficult for individual patients to adequately understand either the specific processes by which decisions are made or the system's future behavior. The black-box problem that pervades AI systems means that only a limited number of specialists can understand such systems sufficiently (30). Furthermore, in systems like MAMLS that engage in continual learning, the scope and impact of data use change over time, and the full picture cannot be grasped in advance (1). Under these conditions, the ideal of "fully informed consent" is difficult to realize in practice.

An anticipated counterargument holds that "complete understanding is not a necessary condition for consent; understanding information that is substantially relevant to decision-making is sufficient" (31). Certainly, informed consent theory does not require complete understanding of technical details. However, in the case of MAMLS, the problem is not limited to the degree of understanding; rather, it lies in the fact that the very object of understanding changes over time. The mode of use understood by the patient at the time of consent may subsequently change through continual learning, and that understanding does not necessarily extend to future uses. Consent in MAMLS is therefore marked by a structural gap between understanding and actual use.

In such circumstances, there is no doubt that efforts to explain technical matters in accessible language and to facilitate patients' understanding are important. However, such "translation" is bound to remain inherently incomplete. A structural asymmetry exists between expert knowledge and ordinary understanding, and it is difficult to overcome this asymmetry entirely. In this respect, consent may not always function as a choice grounded in adequate understanding.

Sparrow and colleagues themselves recognize this point to some degree, noting that IRBs must determine on a case-by-case basis whether consent is required, and that replacing written consent with verbal consent or waiving consent may in some cases be permissible (6). From the perspective of this paper, however, this limitation should not be understood as a problem that can be resolved through procedural adjustments alone; rather, in the context of AI systems such as MAMLS, it calls for a reconsideration of consent as a principle of justification.

Under these circumstances, the character of consent is changing. Conventionally understood as an expression of autonomous choice grounded in understanding, consent in MAMLS increasingly takes the form of an act that presupposes trust in institutions and professionals. Patients do not choose on the basis of having fully grasped the details of data use; rather, they accept such use on the premise that it will be appropriately managed in the future. Consent is thus shifting from "choice based on understanding"

to "trust-based delegation". Biobanks present a structurally analogous case: participants must consent on the basis of trust in the biobank system without knowing in advance the details of future research uses. This structural similarity has already been noted in international discussions, and the accumulated insights from biobank ethics are highly relevant here (32-34).

Importantly, this trust is not limited to trust in individual healthcare providers. Rather, it is directed toward the broader institutional framework that establishes norms for data use and monitors and enforces those norms—regulatory authorities, ethics review bodies, academic institutions, scholarly societies, academic journals, and wider structures of social governance. The focus of ethical justification is thus shifting from the presence or absence of individual consent to the legitimacy of the institutional system as a whole. Such institutional trust requires more than general reassurance. It must be supported by concrete arrangements, including transparency about data use and model updating, accountability mechanisms for harms and biases, continuous oversight after deployment, regular evaluation of model performance across subpopulations, patient and public involvement in governance, and clearly defined roles for regulatory bodies, ethics committees, healthcare institutions, professional societies, and academic journals. These arrangements can help transform trust from a merely psychological attitude into a justified confidence in institutional practices.

Accordingly, while the distinction in consent form—opt-in versus opt-out—remains an important topic of discussion, it does not constitute the core of the problem. This is not to deny the normative significance of opt-in consent: consent that is procedurally valid but substantively thin can still play a symbolic role in protecting patient autonomy and helping to sustain trust in institutions. However, this is not a sufficient condition, and properly structuring consent procedures does not in itself amount to the establishment of institutional trust.

More important is the question of through what processes rules governing data use are formed, and how those processes can acquire public trust. In other words, the ethical challenges surrounding MAMLS need to be repositioned not as a matter of individual choice, but as a problem of building institutional trust. In this respect, the ethics of medical AI extends beyond the conventional research ethics framework centered on consent and connects to a broader theory of social justification—namely, public health ethics and the domain of political philosophy that provides its theoretical foundations. What is at stake there is how to reconcile respect for individual choice with the realization of collective goods, and the resolution of this problem depends to a considerable extent on the construction of a trustworthy institutional framework.

To clarify how the technical and institutional features of MAMLS give rise to ethical challenges and corresponding governance responses, Table 1 summarizes the central structure of the argument.

## 5. Conclusion

This paper has clarified, through two pathways, that in the context of MAMLS the framework of consent has become difficult to sustain as a viable premise. First, the irreversibility inherent in deep learning models substantially undermines withdrawability, an important ancillary right of consent, thereby suggesting that consent may be transformed from a guarantee of ongoing self-determination into something closer to delegation to institutions. Second, the problem of data representativeness and bias shifts the unit of ethical analysis from the individual to the population, bringing to light an "autonomy dilemma" in which efforts within the traditional research ethics framework to respect individual autonomy may paradoxically undermine the effective protection of data providers throughout the population, including their autonomy-related interests.

Under these conditions, the foundation of ethical justification appears to be shifting from the presence or absence of individual consent to trust in the institutional system as a whole. Consent-centered ethics continues to play important normative and legal roles, but under the technological and institutional conditions of MAMLS, it remains open to question whether consent alone can suffice as a principle of justification.

In this respect, while taking the argument of Sparrow and colleagues as a point of departure, this paper seeks to extend its scope one step further. The heart of the problem lies not in the thoroughgoing

implementation of opt-in consent, but in the question of what framework of justification may still retain its validity under technological and institutional conditions in which the consent-based interpretation of respect for persons within the traditional research ethics framework is proving insufficient. This question cannot be adequately examined within the framework of research ethics alone, but must instead be pursued in connection with public health ethics and political philosophy.

Recent developments in the review and amendment of the Act on the Protection of Personal Information in Japan also suggest the practical implications of this problem. The fact that the consent principle itself is being institutionally reconfigured in the context of AI development appears to suggest that the conventional framework, which places consent at the center of ethical justification, is being compelled to reconsider its practical sustainability. These developments may be rooted in a deeper institutional transformation that cannot be adequately captured by the binary of strengthening versus relaxing consent.

Therefore, the tasks ahead are not limited to improving methods for obtaining consent. Rather, the broader question of governance must be confronted: how to institutionally secure the legitimacy of data use and public trust among data providers, who should be involved in this process, and how deliberation and collective decision-making should be organized. In this context, alongside transparency, accountability, and fairness, institutional designs that enable the participation of diverse stakeholders will be essential.

From this perspective, what is being asked in AI medicine is not the technical or procedural question of how consent should be obtained. Rather, under institutional conditions in which consent-centered

**Table 1. Key ethical challenges of Medical Adaptive Machine Learning Systems (MAMLS) and corresponding governance responses**

Technical / institutional feature	Ethical / scientific challenge	Limit of consent-centered ethics	Governance response
Continual learning and data integration	Irreversibility and weakened withdrawability	Withdrawal cannot always be fully implemented after data have influenced model parameters	Transparency about limits of withdrawal; accountability for data use and downstream effects
Model updating over time	Uncertain and evolving future uses	Consent at one time point cannot fully cover later model behavior or new uses	Continuous oversight after deployment; periodic review of model updates and risks
Large-scale data requirements	Selection bias and underrepresentation	Opt-in participation may reduce representativeness and reinforce bias	Fairness monitoring; population-level evaluation; attention to underrepresented groups
Population-level impact	Distribution of risks and benefits across groups	Individual consent alone does not address collective harms	Public health ethics; participatory governance; patient and public involvement
Institutional dependence	Trust-based delegation	Consent alone cannot justify data use when understanding and control are limited	Institutional transparency; regulatory and ethics oversight; clear responsibility structures

approaches to ethical justification are being called into question, we are confronted with a more fundamental set of questions: what do we entrust, to whom, how can that delegation be justified, and how can public trust in healthcare—including medical research—be secured?

**Funding:** This work was supported by AMED under Brain/MINDS 2.0 (Multidisciplinary Frontier Brain and Neuroscience Discoveries), Grant Numbers JP24wm0625001 and JP24wm0625012, and by the Japan Society for the Promotion of Science (JSPS) through a Grant-in-Aid for Scientific Research (A), "A Comprehensive Study on Moral Distress", Grant Number 23H00005.

**Conflict of Interest:** The authors have no conflicts of interest to disclose.

## References

- Vokinger KN, Feuerriegel S, Kesselheim AS. Continual learning in medical devices: FDA's action plan and beyond. *Lancet Digit Health*. 2021; 3:e337-e338.
- Yuan H. Toward real-world deployment of machine learning for health care: External validation, continual monitoring, and randomized clinical trials. *Health Care Sci*. 2024; 3:360-364.
- Bruno P, Quarta A, Calimeri F. Continual learning in medicine: a systematic literature review. *Neural Process Lett*. 2025; 57:Article 2.
- National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (NCPHSBBR). The Belmont report: Ethical principles and guidelines for the protection of human subjects of research. <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html> (accessed May 21, 2026).
- Levine RJ. *Ethics and regulation of clinical research*. 2nd ed. New Haven (CT): Yale University Press; 1986.
- Sparrow R, Hatherley J, Oakley J, Bain C. Should the use of adaptive machine learning systems in medicine be classified as research? *Am J Bioeth*. 2024; 24:58-69.
- Personal Information Protection Commission, Japan. Interim report on considerations for the triennial review of the Act on the Protection of Personal Information. [https://www.ppc.go.jp/files/pdf/240627\\_02\\_houdou\\_betten1.pdf](https://www.ppc.go.jp/files/pdf/240627_02_houdou_betten1.pdf) (accessed April 23, 2026).
- Personal Information Protection Commission, Japan. Results of public comments on the interim report concerning the triennial review of the Act on the Protection of Personal Information. [https://www.ppc.go.jp/files/pdf/240904\\_shiryuu-1-2.pdf](https://www.ppc.go.jp/files/pdf/240904_shiryuu-1-2.pdf) (accessed April 23, 2026). (in Japanese)
- Personal Information Protection Commission, Japan. System Reform Policy under the Triennial Review of the Act on the Protection of Personal Information. [https://www.ppc.go.jp/en/topix/triennial\\_review\\_2026\\_02/](https://www.ppc.go.jp/en/topix/triennial_review_2026_02/) (accessed May 28, 2026).
- Balch JA, Evans BJ, Shickel B, Bihorac A, Upchurch GR Jr, Loftus TJ. The dilemma of consent for AI in healthcare. *Surgery*. 2024; 175:1456-1457.
- Vayena E, Blasimme A, Cohen IG. Machine learning in medicine: Addressing ethical challenges. *PLoS Med*. 2018; 15:e1002689.
- Price WN 2nd, Cohen IG. Privacy in the age of medical big data. *Nat Med*. 2019; 25:37-43.
- Guan H, Bates D, Zhou L. Keeping medical AI healthy and trustworthy: A review of detection and correction methods for system degradation. *IEEE Trans Biomed Eng*. 2025; PP:10.1109/TBME.2025.3642706.
- Topol EJ. High-performance medicine: The convergence of human and artificial intelligence. *Nat Med*. 2019; 25:44-56.
- Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, van der Laak JAWM, van Ginneken B, Sánchez CI. A survey on deep learning in medical image analysis. *Med Image Anal*. 2017; 42:60-88.
- LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015; 521:436-444.
- Bourtole L, Chandrasekaran V, Choquette-Choo CA, Jia H, Travers A, Zhang B, Lie D, Papernot N. Machine unlearning. In: 2021 IEEE Symposium on Security and Privacy (SP). Piscataway (NJ): IEEE; 2021. p. 141-159.
- Xu H, Zhu T, Zhang L, Zhou W, Yu PS. Machine unlearning: A survey. *ACM Computing Surveys*. 2023; 56: Article 9.
- Beauchamp TL, Childress JF. *Principles of biomedical ethics*. 8th ed. New York (NY): Oxford University Press; 2019.
- Manson NC, O'Neill O. *Rethinking informed consent in bioethics*. Cambridge: Cambridge University Press; 2007.
- Miller FG, Wertheimer A, editors. *The ethics of consent: Theory and practice*. New York (NY): Oxford University Press; 2009.
- World Medical Association. *World Medical Association Declaration of Helsinki: Ethical principles for medical research involving human participants*. <https://www.wma.net/policies-post/wma-declaration-of-helsinki/> (accessed April 23, 2026).
- Council for International Organizations of Medical Sciences. *International ethical Guidelines for Health-related research involving humans*. Geneva: Council for International Organizations of Medical Sciences; 2016.
- Wiertz S, Boldt J. Evaluating models of consent in changing health research environments. *Med Health Care Philos*. 2022; 25:269-280.
- Daneshjou R, Vodrahalli K, Novoa RA, et al. Disparities in dermatology AI performance on a diverse, curated clinical image set. *Sci Adv*. 2022; 8:eabq6147.
- Guo LN, Lee MS, Kassamali B, Mita C, Nambudiri VE. Bias in, bias out: Underreporting and underrepresentation of diverse skin types in machine learning research for skin cancer detection-A scoping review. *J Am Acad Dermatol*. 2022; 87:157-159.
- Yamamoto K, Ibuki T, Nakazawa E. The fine balance between complete data integrity in medical adaptive machine learning systems and the protection of research participants. *Am J Bioeth*. 2024; 24:101-103.
- Childress JF, Faden RR, Gaare RD, Gostin LO, Kahn J, Bonnie RJ, Kass NE, Mastroianni AC, Moreno JD, Nieburg P. Public health ethics: Mapping the terrain. *J Law Med Ethics*. 2002; 30:170-178.
- Kass NE. An ethics framework for public health. *Am J Public Health*. 2001; 91:1776-1782.
- Lipton ZC. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*. 2018; 16:31-57.
- Faden RR, Beauchamp TL. A history and theory of

- informed consent. Oxford University Press, New York (NY): Oxford University Press; 1986.
32. Caulfield T, Kaye J. Broad consent in biobanking: Reflections on seemingly insurmountable dilemmas. *Med Law Int.* 2009; 10:85-100.
  33. Steinsbekk KS, Myskja BK, Solberg B. Broad consent versus dynamic consent in biobank research: Is passive participation an ethical problem? *Eur J Hum Genet.* 2013; 21:897-902.
  34. Kaye J, Whitley EA, Lund D, Morrison M, Teare H, Melham K. Dynamic consent: A patient interface for twenty-first century research networks. *Eur J Hum Genet.* 2015; 23:141-146.

----

Received April 28, 2026; Revised May 28, 2026; Accepted June 2, 2026.

Released online in J-STAGE as advance publication June 13, 2026.

*\*Address correspondence to:*

Keiichiro Yamamoto, Department of Clinical Research Management, Center for Clinical Sciences, Japan Institute for Health Security, 1-21-1 Toyama, Shinjuku-ku, Tokyo 162-8655, Japan.

E-mail: yamamoto.kei@jihs.go.jp